arXiv:1512.03127v1 [math.LO] 10 Dec 2015

FLEXIBLE CONSTRAINT SATISFIABILITY AND A PROBLEM IN SEMIGROUP THEORY

MARCEL JACKSON

Abstract. We examine some flexible notions of constraint satisfaction, observing some relationships between model theoretic notions of universal Horn class membership and robust satisfiability. We show the NP-completeness of 2-robust monotone 1-in-3 3SAT in order to give very small examples of finite algebras with NP-hard variety membership problem. In particular we give a 3-element algebra with this property, and solve a widely stated problem by showing that the 6-element Brandt monoid has NP-hard variety membership problem. These are the smallest possible sizes for a general algebra and a semigroup to exhibit NP-hardness for the membership problem of finite algebras in finitely generated varieties.

In a number of computational situations, the task is not to find if there is a single solution to an instance, but rather to find whether there is a sufficiently broad family of solutions to witness various separation conditions on variables. The following example is illustrative of idea.

Recall that a 3-*colouring* of a graph $\mathbb{G}$ is a homomorphism from $\mathbb{G}$ to the complete graph $\mathbb{K}_3$. Equivalently, there is a map to the set $\{0, 1, 2\}$ such that adjacent vertices have different colours. Graph colouring places no restriction on what is to happen to *non-adjacent* vertices. What happens then if we ask that a graph not simply be 3-colourable, but 3-colourable in enough ways that there is no restriction to how any pair of nonadjacent vertices may be coloured: can they be coloured the same by some valid 3-colouring, and can they be coloured differently? This interesting combinatorial question can be rephrased in the following equally interesting way: is it true that every valid colouring of two vertices extends to a full 3-colouring?

In the present article we look at a range of constraint problems where the question of satisfaction is replaced by the question as to whether there is a sufficiently flexible array of solutions to universally achieve various separation and identification properties on variables. We group this general notion loosely under the name *flexible satisfaction*, with specific kinds of flexible satisfaction to be given more precise definition. The idea appears natural enough in its own right, and so it is perhaps not surprising that it emerges from several different investigations of study. For example, there has been extensive exploration of the emergence of computationally challenging instances in randomly generated instances of constraint problems; see Hayes [16] or Monasson et al. [32] for example. Such work points toward the emergence of hidden constraints being intimately related to computational difficulty; see Beacham and Culberson [3] or Culberson and Gent [8] for example. These hidden

2010 *Mathematics Subject Classification.* Primary: 68Q17, 20M07; Secondary: 03C05, 08B99, 08C15.

The author was supported by ARC Discovery Project DP1094578 and ARC Future Fellowship FT120100666.

constraints correspond exactly to non-constrained tuples being preserved as a constraint under any possible solution, or being forced to fall outside of any relation under any solution (cf. non-adjacent vertices being forced to be differently coloured or forced to be like-coloured, respectively). As a second example, the resolution of problems relating to minimal constraint problems (in the sense of Montenari [33]), led Gottlob [14] to *supersymmetric* SAT instances, which concerns precisely the problem of extendability of arbitrary partial solutions (on a bounded number of variables) to full solutions. This is developed further in Abramsky, Gottlob and Kolaitis [2] in the guise of *robust satisfiability*, where the flexible graph 3-colouring problem described above is shown to be NP-complete, and related notions are tied to a problem in quantum mechanics. A third point of interest lies in the fundamental relationship between flexible satisfaction and universal Horn classes: the flexibly 3-colourable graphs (in the sense described above) are precisely those in the universal Horn class generated by $\mathbb{K}_3$; see Trotta [37] for example.

A final point of interest arises somewhat more coincidentally. Intuitively, constraint networks that can be satisfied in very flexible ways lie at some extreme antipodal point compared to those that are not even satisfiable at all. Yet in the course of proving NP-completeness, it seems necessary to use a string of reductions which leads to an interesting stretching of the landscape of instances. It is found that there is no polynomial time boundary separating the ostensibly extreme subclasses *not satisfiable* and *flexibly satisfiable*. This can already be seen in the results of Abramsky, Gottlob and Kolaitis [2], where after reduction, one arrives at a graph that is either not 3-colourable at all, or which is flexibly 3-colourable. It is also built into the definition of supersymmetric SAT in [14]. This extreme separation of YES instances from NO instances gives the results extra applicability, which we make use of in the present article. Aside from wider applicability, the issue is also of interest from a purely computational complexity perspective, as it is at the very least quite reminiscent of the currently open problem as to whether or not there are two disjoint NP languages that are inseparable by a polynomial time boundary (see [10, 11, 15], for example), and is a natural example of a promise problem (see the survey by Goldreich [12] for example).

A key motivation for the present article comes from computational problems arising in the theory of varieties of algebras. Recall that the *variety* generated by an algebra $\mathbf{A}$ is the class of algebras $\mathbf{B}$ of the same signature as $\mathbf{A}$ and which satisfy all identities true on $\mathbf{A}$. Equivalently, $\mathbf{B}$ is in the variety generated by $\mathbf{A}$ if and only if it is a homomorphic image of a subalgebra of a direct power of $\mathbf{A}$. The pseudovariety of $\mathbf{A}$ is the finite part of the variety of $\mathbf{A}$, though in general a *pseudovariety* will be any class of finite algebras of the same signature and which is closed under taking homomorphic images, subsemigroups and finitary direct products. Not every pseudovariety is the finite part of a variety, so the pseudovariety setting is a more encompassing setting for considering finite membership problems of finite algebras.

Bergman and Slutzki [5] gave a 2-exptime algorithm for deciding membership of $\mathbf{B}$ in the pseudovariety of $\mathbf{A}$ (even if both algebras are allowed to vary), while Kozik [27] showed that even if $\mathbf{A}$ is fixed, this is best possible in general: there is a finite algebra with 2-exptime complete finite membership problem for its pseudovariety. Other examples with nontractable finite membership problem (subject to usual complexity theoretic assumptions such as $\mathtt{P} \neq \mathtt{NP}$) include the original

example, due to Szekely [36], of a 7-element (or 6-element [19]) algebra with NP-complete membership problem, semigroups of the author and McKenzie [19] with NP-hard membership problem (the smallest has 55 elements, though for a monoid one needs 56 elements), a finite algebra due to Kozik with pspace-complete [26] finite membership problem for its pseudovariety, and a family of finite semigroups recently discovered Klíma, Kunc and Polák [25], each with co-NP-complete finite membership problem for their pseudovariety (the smallest is a monoid of size 42).

All of these examples are somewhat ad hoc. We are able to use the hardness of various flexible satisfaction ideas—and in particular, the extreme separation properties that arise—to give an interval in the lattice of semigroup pseudovarieties in which every pseudovariety has NP-hard finite membership problem. At the base of this interval is the pseudovariety generated by the ubiquitous six element Brandt monoid $\mathbf{B}_2^1$. The particular case of $\mathbf{B}_2^1$ solves—subject to the assumption $\mathtt{P} \neq \mathtt{NP}$—Problem 4 of Almeida [1, p. 441] and Problem 3.11 of Kharlampovich and Sapir [23]; see also Volkov, Gol′dberg and Kublanovskiĭ [38, p. 849] (English version). Also, because every semigroup of order less than 6 has a finite identity basis (hence tractable membership problem for its pseudovariety), it shows that the smallest size of generator for a pseudovariety with computationally hard membership is 6. For general algebras, we show that there is a 3-element algebra with NP-complete membership problem for its pseudovariety variety, also the smallest possible size; we give a 4-element groupoid (that is, an algebra with a single binary operation) with the same property.

0.1. **Results and structure.** The structure of the article is as follows. In Part 1 we give preliminary discussion of basic concepts relating to constraint satisfaction problems, and relate various flexible notions of satisfaction to model theoretic classes such as universal Horn classes. Results here are mostly basic observations, but nevertheless of some interest. Theorem 2.1 identifies unfrozenness of relations to the structure of Horn clauses defining the class; several examples are given, including Example 2.2 which shows that the not-necessarily induced subgraphs of direct powers of the complete graph $\mathbb{K}_3$ are those 3-colourable graphs in which nonadjacent vertices can be distinctly coloured. Proposition 3.2 gives a logspace relationship between CSP complexity and the complexity of deciding unfrozenness and frozenness of relations in the case of core structures; the relationship is shown to fail for non-core structures in Example 3.3.

Part 2 contains the most technical sections of the article. We build off the work of Gottlob [14] to establish a series of results about robust satisfiability, culminating in the NP-completeness of 2-robust monotone 1-in-3 3SAT; see Theorem 6.1. The result also shows that the universal Horn class generated by the 2-element relational template for monotone 1-in-3 3SAT (which consists of the single ternary relation $\{(1,0,0),(0,1,0),(0,0,1)\}$ on the set $\{0,1\}$) has NP-complete membership problem. As alluded to above, every boundary separating the non-satisfiable instances from those that lie in this universal Horn class has NP-hard membership problem. Along the way, we reprove the NP-completeness of deciding flexible 3-colourability; this is Theorem 5.1. We prove this for graphs in which every edge lies within a triangle and show also that the flexibility of colouring extends slightly further to a kind of edge flexibility: for any two edges, we show that aside from a few obvious special cases where the range of colours are limited, otherwise we can always use 2 colours in total amongst the two edges, and also 3 colours; see Proposition 5.2. The main

difficulties relate to the triangulation of edges, however this is the detail we require for a reduction into monotone 1-in-3 3SAT.

Part 3 contains applications of the results in Part 2 to hardness of finite membership problems for pseudovarieties. In Section 7 we consider the pseudovariety $\mathcal{V}(\mathbf{B}_2^1)$ generated by the six-element Brandt monoid $\mathbf{B}_2^1$ and show in Theorem 7.3 that any pseudovariety containing $\mathcal{V}(\mathbf{B}_2^1)$ and contained within the join of $\mathcal{V}(\mathbf{B}_2^1)$ with $\mathbb{L}\mathbf{DS}$ has NP-hard finite membership problem (with respect to logspace reductions). Note that $\mathbb{L}\mathbf{DS}$ is the largest pseudovariety omitting $\mathbf{B}_2^1$. As a corollary we deduce that adding an identity element to any finite completely 0-simple semigroup whose sandwich matrix is block-diagonal produces a semigroup (or monoid) whose pseudovariety has NP-hard finite membership problem; Corollary 7.4. In Section 8 we deduce some easier consequences of the flexible satisfaction results, giving a 3-element algebra and 4-element groupoid with NP-complete and NP-hard finite membership problem (respectively) for their pseudovarieties; see Corollary 8.1.

Part 1. Preliminaries: constraints, satisfaction and separation

1. Constraint Satisfaction Problems.

In this article we consider constraint problems over fixed finite relational templates only. A *template* $\mathbb{A}$ consists of a finite set A (the *domain*) endowed with a finite family $R_1^{\mathbb{A}}, \ldots, R_n^{\mathbb{A}}$ of relations on A, each of some finite arity. The family of symbols $R_1, \ldots, R_n$ and associated arities is the *vocabulary*, or *signature* of $\mathbb{A}$, so that we can equivalently think of the template $\mathbb{A}$ as a relational structure $\mathbb{A} = \langle A; R_1^{\mathbb{A}}, \ldots, R_n^{\mathbb{A}} \rangle$.

Two key examples of interest in the present article are the template for graph 3-colourability, which consists of the simple graph

$$\mathbb{K}_3 = \langle \{0,1,2\}; \neq \rangle$$

which has a single relation, of arity 2, and the two-element template for monotone 1-in-3 3SAT

$$2 = \langle \{0,1\}; \{(0,0,1),(0,1,0),(1,0,0)\} \rangle$$

which has a single relation, of arity 3.

The *constraint satisfaction problem* $\mathrm{CSP}(\mathbb{A})$ over $\mathbb{A}$ is the computational problem that takes as an instance a set of variables $x_1, \ldots, x_m$ and a set of constrained tuples: expressions of the form $(x_{i_1}, \ldots, x_{i_k}) \in R$, where R is a symbol of arity k from the vocabulary of $\mathbb{A}$. Such an instance is a YES instance of $\mathrm{CSP}(\mathbb{A})$ if there is an assignment $\phi : \{x_1, \ldots, x_m\} \to A$ satisfying all of the constraints of the instance: for each constraint $(x_{i_1}, \ldots, x_{i_k}) \in R_j$, we must have $(\phi(x_{i_1}), \ldots, \phi(x_{i_k})) \in R_j^{\mathbb{A}}$. Otherwise the instance is a NO instance of $\mathrm{CSP}(\mathbb{A})$.

An algebraic/model-theoretic perspective reveals that $\mathrm{CSP}(\mathbb{A})$ is nothing other than the homomorphism problem for $\mathbb{A}$. To see this, we may consider the variables in a CSP instance I as the universe of some relational structure $\mathbb{B}$ of the same signature as $\mathbb{A}$. For each symbol R in the relational signature of $\mathbb{A}$, the relation $R^{\mathbb{B}}$ consists of the set of all tuples that were constrained to be R-related in I:

$$R^{\mathbb{B}} := \{(x_{i_1}, \ldots, x_{i_k}) \mid (x_{i_1}, \ldots, x_{i_k}) \in R \text{ is a constraint in } I\}.$$

Then the notion of $\nu : \{x_1, \ldots, x_m\} \to A$ satisfying the constraints of I coincides with the definition of $\nu : \mathbb{B} \to \mathbb{A}$ being a homomorphism. Conversely, any finite relational structure $\mathbb{B}$ of the same type as $\mathbb{A}$ gives a CSP instance by calling the

elements of the universe of $\mathbb{B}$ "variables" and taking the list of elements of each fundamental relation of $\mathbb{B}$ as the list of constraints.

2. Separation, identification and unfrozen relations

Let $\mathcal{R}$ be a relational signature and $R \in \mathcal{R} \cup \{=\}$. For a template $\mathbb{A}$ of signature $\mathcal{R}$, we will say that an $\mathcal{R}$-structure $\mathbb{B}$ satisfies the *R-separation condition* with respect to $\mathbb{A}$ if for each $(b_1, \ldots, b_k) \notin R^{\mathbb{B}}$, there is a satisfying assignment ν (that is, a homomorphism $\nu : \mathbb{B} \to \mathbb{A}$) with $(\nu(b_1), \ldots, \nu(b_k)) \notin R^{\mathbb{A}}$, where of course $=$ to be interpreted as equality on both $\mathbb{A}$ and $\mathbb{B}$. The following notion is somewhat dual to R-separation: when $(b_1, \ldots, b_k) \notin R^{\mathbb{B}}$ there is a satisfying assignment ν' with $(\nu(b_1), \ldots, \nu(b_k)) \in R^{\mathbb{A}}$. We refer to this property as the *R-identification condition* with respect to $\mathbb{A}$.

Following Beacham [4], we may say that a tuple $(b_1, \ldots, b_k) \notin R^{\mathbb{B}}$ is *frozen-in* (with respect to $\mathbb{A}$) if the separation condition fails at $(b_1, \ldots, b_k) \notin R^{\mathbb{B}}$. Similarly, the tuple $(b_1, \ldots, b_k) \notin R^{\mathbb{B}}$ is *frozen-out* if the R-identification property fails. We say that the relation $R^{\mathbb{B}}$ is *unfrozen-in* (with respect to $\mathbb{A}$) if no tuple outside of $R^{\mathbb{B}}$ is frozen in. Dually, we may speak of unfrozen-out relations on $\mathbb{B}$ (with respect to $\mathbb{A}$). As we now see, these frozen properties are closely related to familiar concepts in model theory and universal algebra. The following definitions can be found Burris and Sankappanavar [7] for example; a more detailed treatment is Gorbunov [13].

A *Horn clause* in a first order language (with equality) is a disjunction

$$\phi_1 \vee \phi_2 \vee \cdots \vee \phi_k \tag{1}$$

in which each of $\phi_1, \ldots, \phi_k$ is either atomic or negated-atomic, but *at most one* disjunct ϕ_i is not negated. A first order sentence in this language is a *universal Horn sentence* if it is a universally quantified Horn clause. (Technically one might take universally quantified conjunctions of Horn clauses, but this is equivalent to a finite set of sentences of the form we have described.)

When *exactly one* disjunct is not negated (say, ϕ_k), the Horn clause (1) can be written as the implication

$$(\neg\phi_1 \ \& \ \cdots \ \& \ \neg\phi_{k-1}) \to \phi_k \tag{2}$$

(of course as $\phi_1, \ldots, \phi_{k-1}$ were already negated by assumption, the premise of (2) is really just a conjunction of atomic formulæ). Universally quantified expressions of this form are usually called *quasi-equations* (or quasi-identities). Universal Horn sentences in which *all* disjuncts are negated are often called *anti-identities*, though these are also easily seen to be just the negations of primitive-positive sentences (existentially quantified conjunctions of atomic formulæ).

The following theorem is a basic synthesis of a number of commonly encountered conditions. When $k = 0$ (meaning $\{R_1, \ldots, R_k\} = \varnothing$), then Theorem 2.1 gives the usual syntactic characterisation for the antivariety of a finite structure (or, when restricted to finite $\mathbb{B}$). When $\{R_1, \ldots, R_k\}$ is $\mathcal{R} \cup \{=\}$, the theorem gives a very widely used series of equivalent conditions for membership in the universal Horn class generated by a finite structure. When $k = 1$ and R_1 is the relation $=$, then (2) $\Leftrightarrow$ (3) of the theorem is essentially Corollary 8 of Stronkowski [35]. (When $\mathbb{A}$ is infinite, or if we consider an infinite set of structures instead of a single $\mathbb{A}$, then one needs to also account for ultraproduct closure, but there is an obvious extension of statements 1 and 3 that holds, with standard adjustments to the proof.)

Theorem 2.1. *Let $\mathbb{A}$ be a finite relational structure and $\mathbb{B}$ a relational structure of the same finite signature $\mathcal{R}$. Let $\{R_1, \ldots, R_k\}$ be a subset of $\mathcal{R} \cup \{=\}$. The following are equivalent.*

(1) $\mathbb{B} \in \mathrm{CSP}(\mathbb{A})$ and satisfies the separation condition for each of $R_1, \ldots, R_k \in \mathcal{R} \cup \{=\}$ (equivalently, each of $R_1, \ldots, R_k$ is unfrozen-in).
(2) $\mathbb{B}$ satisfies all universal Horn sentences satisfied by $\mathbb{A}$ for which any non-negated disjunct involves relations only from $R_1, \ldots, R_k$.
(3) There is a homomorphism from $\mathbb{B}$ into a finite nonzero direct power of $\mathbb{A}$ that preserves $\{\neg R_1, \ldots, \neg R_k\}$.

Proof. (1)⇒(3). Assume (1) and let $\hom(\mathbb{B}, \mathbb{A})$ be the family of all homomorphisms from $\mathbb{B}$ into $\mathbb{A}$; this is nonempty because $\mathbb{B} \in \mathrm{CSP}(\mathbb{A})$. There is a natural homomorphism from $\mathbb{B}$ into $\mathbb{A}^{\hom(\mathbb{B},\mathbb{A})}$ given by identifying each b with the function $e_b : \hom(\mathbb{B}, \mathbb{A}) \to \mathbb{A}$ defined by $e_b(\phi) = \phi(b)$. This is a homomorphism, but also it preserves $\{\neg R_1, \ldots, \neg R_k\}$: for each $R \in \{R_1, \ldots, R_k\}$ and $(b_1, \ldots, b_\ell) \notin R^{\mathbb{B}}$, the separation condition of (1) ensures that there is $\phi : \mathbb{B} \to \mathbb{A}$ with $(\phi(b_1), \ldots, \phi(b_\ell)) \notin R^{\mathbb{B}}$, so that $(e_{b_1}, \ldots, e_{b_\ell})(\phi) = (\phi(b_1), \ldots, \phi(b_\ell)) \notin R^{\mathbb{A}}$. Thus (3) holds.

(3)⇒(1). Assume (3) holds. Let $\mathbb{B}'$ be the image of $\mathbb{B}$ under the assumed homomorphism. Then whenever $(b_1, \ldots, b_\ell) \notin R^{\mathbb{B}}$ for some $R \in \{R_1, \ldots, R_k\}$ we also have $(b_1, \ldots, b_\ell) \notin R^{\mathbb{B}'}$, and projecting onto a coordinate i in which $(b_1(i), \ldots, b_\ell(i)) \notin R^{\mathbb{A}}$ gives the desired homomorphism from $\mathbb{B}$ into $\mathbb{A}$ witnessing the R-separation condition at the (arbitrary) tuple $(b_1, \ldots, b_\ell) \notin R^{\mathbb{B}}$. This shows that (1) holds.

(2)⇒(1). Assume that (1) fails. Let $\mathrm{diag}^+(\mathbb{B})$ denote positive diagram of $\mathbb{B}$. If $\mathbb{B} \notin \mathrm{CSP}(\mathbb{A})$, then $\mathbb{A}$ satisfies $\neg\, \mathrm{diag}^+(\mathbb{B})$, which is equivalent to an anti-identity, and which trivially fails on $\mathbb{B}$. Now assume that $\mathbb{B} \in \mathrm{CSP}(\mathbb{A})$. It follows then that for some $R \in \{R_1, \ldots, R_k\}$, the R-separation condition fails at a tuple $(b_1, \ldots, b_\ell) \in B^\ell \backslash R^{\mathbb{B}}$. Then $\mathbb{A}$ satisfies the quasi-identity $\mathrm{diag}^+(\mathbb{B}) \to (b_1, \ldots, b_\ell) \in R$, while again $\mathbb{B}$ trivially fails this. Hence (2) fails.

(1)⇒(2). Finally, assume that (2) fails for some universal Horn sentence ψ. Let $x_1, \ldots, x_n$ be the variables appearing in ψ, and let $b_1, \ldots, b_n$ be elements of B such that $\mathbb{B}$ fails ψ at the substitution $x_i \mapsto b_i$. If ψ is an anti-identity, then $\mathbb{B} \notin \mathrm{CSP}(\mathbb{A})$. Now assume that one disjunct in ψ is an atomic formula $(x_{i_1}, \ldots, x_{i_\ell}) \in R$ (with R in $\{R_1, \ldots, R_k\}$), so that it can be written as a quasi-identity, with conclusion $(x_{i_1}, \ldots, x_{i_\ell}) \in R$. So, as ψ fails at $b_1, \ldots, b_n$, we have that $(b_{i_1}, \ldots, b_{i_\ell}) \notin R^{\mathbb{B}}$. Let $\phi : \mathbb{B} \to \mathbb{A}$ be an arbitrary homomorphism. Then ϕ yields a satisfying interpretation of the premise of ψ in $\mathbb{A}$ by giving each x_i the value $\phi(b_i)$. As $\mathbb{A} \models \psi$, we have that $(\phi(b_{i_1}), \ldots, \phi(b_{i_\ell})) \in R^{\mathbb{A}}$. Hence the R-separation condition fails, as required. □

In the strongest case of Theorem 2.1—when $\{R_1, \ldots, R_k\}$ is $\mathcal{R} \cup \{=\}$—part (3) states that $\mathbb{B}$ is isomorphic to an induced substructure of a direct power (over nonempty index set) of $\mathbb{A}$. The class of induced substructures of direct powers of $\mathbb{A}$ is denoted by $\mathsf{SP}(\mathbb{A})$ (where S is the class operator returning isomorphic copies of induced substructures, and P is the class operator returning nonempty-index set direct powers of $\mathbb{A}$). The symbol R is used in Stronkowski [35] to denote the class operator of taking not-necessarily-induced substructures. The following example then concerns membership in the class $\mathsf{RP}(\mathbb{K}_3)$.

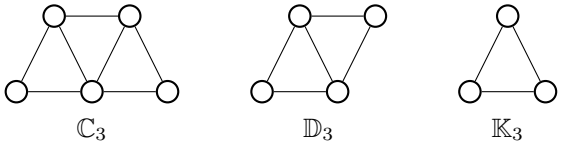

FIGURE 1. Three finite graphs generating universal Horn classes with various stages of unfrozenness relative to homomorphisms into $\mathbb{K}_3$; see Example 2.4.

Example 2.2. *The class of graphs arising a subgraphs (not necessarily induced) of direct powers of the complete graph $\mathbb{K}_3$ is exactly the class of 3-colourable digraphs satisfying the separation condition for $=$. Equivalently, a graph is a subgraph of a nonempty direct power of $\mathbb{K}_3$ if and only if it is 3-colourable, and every pair of nonadjacent vertices can be coloured distinctly.*

Proof. Technically, Theorem 2.1 would apply to directed graphs, but the restriction to symmetric edge relation is easy to see. Let R_1 be the relation of equality. Then by Theorem 2.1, $=$ is unfrozen-in (with respect to $\mathbb{K}_3$) for a graph $\mathbb{G}$ if and only if there is a homomorphism from $\mathbb{G}$ into a nonempty power of $\mathbb{K}_3$, which preserves $\neq$. In other words, 3-colourable, and every pair of nonadjacent vertices can be coloured distinctly if and only if there is an injective homomorphism from $\mathbb{G}$ into a nonempty power of $\mathbb{K}_3$. This last property is equivalent to being isomorphic to a (not-necessarily induced) subgraph of a power of $\mathbb{K}_3$. □

Remark 2.3. In Jackson and Trotta [21, Lemma 6] it is shown that for every finite template $\mathbf{T}$ of finite signature, there is a finite template $\mathbf{T}^\sharp$ of the same signature and with $\mathrm{CSP}(\mathbf{T}) = \mathrm{CSP}(\mathbf{T}^\sharp) = \mathsf{SP}(\mathbf{T}^\sharp)$. In other words, there is always an equivalent template $\mathbf{T}^\sharp$ relative to which all YES instances of $\mathrm{CSP}(\mathbf{T}) = \mathrm{CSP}(\mathbf{T}^\sharp)$ have every relation (and equality) unfrozen-in.

The following examples are used later and also demonstrate structures achieving the "defrosting" of relations in the graph $\mathbb{K}_3$.

Example 2.4. *The examples concern the three graphs shown in Figure 1.*

(1) 3-colourability. The class of 3-colourable graphs coincides with class of (isomorphic copies of) induced subgraphs of direct powers of the graph $\mathbb{C}_3$.

(2) Edge-unfrozen-in 3-colourability. The class of graphs that satisfy, relative to $\mathbb{K}_3$, the separation property for their edge relation coincides with the class of isomorphic copies of induced subgraphs of direct powers of the graph $\mathbb{D}_3$.

(3) Edge- and equality-unfrozen-in 3-colourability. The class of graphs that satisfy, relative to $\mathbb{K}_3$, the separation property for both their edge relation and equality coincides with the class of isomorphic copies of induced subgraphs of direct powers of $\mathbb{K}_3$. This class coincides with the class of graphs that, relative to $\mathbb{K}_3$, satisfy the separation property the identification property for their edge relation.

Proof. Part (1) is due to Nešetřil and Pultr [34], though there is a similar construction also in Wheeler [39]. Part (3) can be found in Trotta [37, Lemma 4.1]. There are obvious extensions to k-colourability given in each of these articles. For part (2), let $\mathbb{F}$ denote the graph determined by the given adjacency matrix, with

vertices $0, 1, 2, 3$, in the given order. Now let $\mathbb{G}$ be a graph whose edge relation is unfrozen-in (relative to $\mathbb{K}_3$). Then the edge relation remains unfrozen-in with respect to $\mathbb{F}$ because $\mathbb{K}_3$ is an induced subgraph of $\mathbb{F}$. We claim that equality is also unfrozen-in with respect to $\mathbb{F}$. Indeed, if $u \neq v$ are distinct vertices from $\mathbb{G}$, take any 3-colouring $\nu : \mathbb{G} \to \mathbb{K}_3$ of $\mathbb{G}$. If $\nu(u) \neq \nu(v)$ we are done because $\mathbb{K}_3$ is an induced subgraph of $\mathbb{F}$. Otherwise, if $\nu(u) = \nu(v)$, then we can assume without loss of generality that $\nu(u) = \nu(v) = 0$. Define a new map from $\mathbb{G}$ into $\mathbb{F}$ by $\nu'(w) := \nu(w)$ for all vertices w, except for v where we define $\nu'(v) = 3$. This is a homomorphism $\nu'(u) \neq \nu'(v)$. So, by Theorem 2.1 it follows that $\mathbb{G}$ is isomorphic to an induced subgraph of a power of $\mathbb{F}$.

Conversely, any induced subgraph $\mathbb{G}$ of a power of $\mathbb{F}$ has unfrozen edges with respect to $\mathbb{K}_3$. To see this, let u and v be such that $\{u, v\}$ is not an edge of $\mathbb{G}$. Then there is a homomorphism $\nu : \mathbb{G} \to \mathbb{F}$ with $\{\nu(u), \nu(v)\}$ not an edge of $\mathbb{F}$. Follow this by the retraction of $\mathbb{F}$ onto the induced subgraph on vertices $\{0, 1, 2\}$, which is isomorphic to $\mathbb{K}_3$. This maps nonedges of $\mathbb{F}$ to nonedges of $\mathbb{K}_3$. Thus we have a 3-colouring of $\mathbb{G}$ in which the nonedge $\{u, v\}$ is mapped to a nonedge of $\mathbb{K}_3$. □

Deciding membership in the class of edge-unfrozen-out 3-colourable graphs was shown to be NP-complete by Beacham [4] (edge unfrozen-out coincides with equality unfrozen-in for $\mathbb{K}_3$). Deciding membership in the class of graphs that are both edge and equality unfrozen for $\mathbb{K}_3$ was shown to be NP-complete by Abramsky, Gottlob and Kolaitis [2]. This property coincides with "2-robust 3-colourability", which we now describe.

3. Robust satisfaction

The various separation and identification conditions have a close relationship with "robust satisfiability" in the sense of Abramsky, Gottlob and Kolaitis [2]. A *partial assignment* from $\mathbb{B}$ to $\mathbb{A}$ is a map ϕ from some subset S of the underlying universe B into $\mathbb{A}$. The notion of robust satisfiability concerns the ability to extend partial assignments to full satisfying assignment. For this it is necessary that ϕ satisfy any existing constraints on S, but also any additionally implicit constraints imposed by the basic constraints on S. In Abramsky, Gottlob and Kolaitis [2], a partial assignment is said to be *locally compatible* if it preserves all projections of constraints, in addition to the constraints themselves. Then the structure $\mathbb{B}$ is *k-robustly satisfiable* if every locally compatible partial assignment from a k element subset of B into $\mathbb{A}$ extends to a satisfying assignment of $\mathbb{B}$ in $\mathbb{A}$.

As an example, consider the template for monotone 1-in-3 3SAT, which consists of the domain $\{0, 1\}$ and the single ternary relation $R = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$. The projection of R to any two of its three coordinates is the binary relation $R' := \{(0, 0), (0, 1), (1, 0)\}$. A locally compatible partial truth assignment on two variables x, y must preserve R' as well as R. In particular if $(x, y) \in R'$ (meaning there is some z such that $(x, y, z) \in R$ is a constraint) then a partial truth assignment would fail to be local compatible if it assigned both x and y the value 1.

There are natural variations of this notion of local compatibility. Recall that a relation is *primitive-positive definable* (abbreviated to *pp-definable*) from a set of relations $\{R_1, \ldots, R_k\}$ if it is the solution set to an existentially quantified formula built from the relations $R_1, \ldots, R_k$ and conjunction. It is easily verified that pp-definable relations must be preserved by any homomorphism preserving $\{R_1, \ldots, R_k\}$. The notion of local compatibility is an example of a pp-definition,

as the projection of a k-ary relation R to some subset $\{i_1, \ldots, i_\ell\} \subseteq \{1, \ldots, k\}$ is the relation $\{(x_{i_1}, \ldots, x_{i_\ell}) \mid \exists x_{j_1} \ldots \exists x_{j_{k-\ell}}\ (x_1, \ldots, x_k) \in R\}$, where $i_1 < \cdots < i_\ell$, $j_1 < \cdots < j_{k-\ell}$ and $\{j_1, \ldots, j_{k-\ell}\} = \{1, \ldots, k\} \backslash \{i_1, \ldots, i_\ell\}$. The following lemma is probably folklore.

Lemma 3.1. *Let $\mathbb{B}$ be an instance of* CSP($\mathbb{A}$)*, and $\nu : S \to A$ be a partial assignment from some subset S of B. Then ν extends to a homomorphism if and only if ν preserves the restriction to S of all relations primitive-positive definable from the fundamental relations of $\mathbb{B}$.*

Proof. Certainly every homomorphism preserves primitive-positive definable relations. So it suffices to show that if ν preserves every pp-definable relation then it extends to a homomorphism. Let $s_1, \ldots, s_\ell$ be an enumeration of the elements of S. Now let $\delta(s_1, \ldots, s_\ell)$ denote the positive diagram of $\mathbb{B}$ (each element of B is a variable) but with all variables formed from $B \backslash S$ existentially quantified and the variables $s_1, \ldots, s_\ell$ formed from S left unquantified. Let R denote the ℓ-ary relation that is pp-defined by $\delta(s_1, \ldots, s_\ell)$; obviously $(s_1, \ldots, s_\ell)$ is a tuple in S that is in this relation on $\mathbb{B}$, and hence on the substructure on S. Assume then that ν preserves the restriction of R to S. So in particular, $(\nu(s_1), \ldots, \nu(s_\ell)) \in R^{\mathbb{A}}$. Now each $b \in B \backslash S$ appears as existentially quantified in δ, and hence there is an interpretation of each $b \in B \backslash S$ in A such that the atomic diagram of $\mathbb{B}$ holds. This means there is a homomorphism from $\mathbb{B}$ into $\mathbb{A}$ that agrees with ν on S. □

Thus, requiring that *all* pp-definable relations be preserved is the same as requiring that ν extend to a homomorphism, which is hardly local. By restricting the number of quantified variables in a pp-formula though, we obtain a natural hierarchy of local compatibility conditions. We could say for example that a relation is *k-local* if it can be defined from the fundamental relations by way of a pp-formula with at most k existentially quantified variables.

Consider for example the relational template for monotone NAE3SAT, which consists of the domain $\{0, 1\}$ with the one ternary relation R given on $\{0, 1\}$ by $\{(0, 0, 1), (0, 1, 0), (1, 0, 0), (1, 1, 0), (1, 0, 1), (0, 1, 1)\}$. Consider the instance on variables v, w, x, y, z with $(v, w, x) \in R$ and $(x, y, z) \in R$. The partial assignment $\nu : \{v, w, y, z\} \to \{0, 1\}$ given by $\nu(v) = \nu(w) = 1$ and $\nu(y) = \nu(z) = 0$ preserves all projections of the relation R, however it fails to extend to a satisfying truth assignment because it fails the pp-definable 4-ary relation $\{(a, b, d, e) \mid \exists c\ (a, b, c) \in R\ \ \&\ \ (c, d, e) \in R\}$, under which (v, w, y, z) is related and which on $\{0, 1\}$ avoids the tuple $(1, 1, 0, 0)$ as well as $(0, 0, 1, 1)$. This example demonstrates that it is almost impossible for a monotone NAE3SAT instance to be 4-robustly satisfiable: it would require that all clauses either do not overlap at all or overlap by two variables. However, asking for preservation of 1-local relations might allow for a much more powerful concept of 4-robust NAE-satisfiability in NAE3SAT.

Recall that a relational structure is a *core* if every endomorphism is an automorphism. Every finite relational structure $\mathbb{A}$ retracts onto a core substructure $\mathbb{A}^\flat$, so that CSP($\mathbb{A}$) is identical to CSP($\mathbb{A}^\flat$). Moreover, CSP($\mathbb{A}^\flat$) is logspace equivalent to the CSP over $\mathbb{A}^\flat$ with all singleton unary relations adjoined [6, 28]. The following is a variation of this logspace equivalence.

Proposition 3.2. *Let $\mathbb{A} = \langle A; R_1, \ldots, R_k \rangle$ be a finite relational structure of finite signature. If $\mathbb{A}$ is a core then for each of the following decision problems there is a logspace Turing reduction from the problem to* CSP($\mathbb{A}$):

(1) *deciding k-robust satisfiability of a structure $\mathbb{B}$ relative to $\mathbb{A}$;*
(2) *deciding of a given structure $\mathbb{B}$ if any given combination of the R-separability and R-identification properties hold, for any subset of relations from $\{R_1, \ldots, R_k\} \cup \{=\}$.*

Proof. The algorithm will consult a dictionary consisting of the complements of each fundamental relation on $\mathbb{A}$: for each fundamental relation R or arity k, construct $S_i := A^k \backslash R^{\mathbb{A}}$. These are fixed, so not part of the logspace reduction. Let $\mathbb{B}$ be an instance of one of the above problems.

To decide k-robustness, we will successively test each locally compatible partial assignment on k points as follows. Given a partial assignment ϕ on k elements of B, first test in logspace that it is locally compatible (because k is a fixed number this can be done in logspace). If not, discard and construct the next possible partial assignment. If ϕ is locally compatible, then write a new structure $\mathbb{B}_\phi$ to the query tape for a CSP($\mathbb{A}$)-oracle as follows. The structure $\mathbb{B}_\phi$ is simply $\mathbb{B}$ with a copy of $\mathbb{A}$ adjoined but with each b_i identified with $\phi(b_i)$. We claim that $\mathbb{B}_\phi \in \mathrm{CSP}(\mathbb{A})$ if and only if ϕ extends to a homomorphism from $\mathbb{B}$ to $\mathbb{A}$. Clearly if ϕ extends to a full assignment of $\mathbb{B}$ into $\mathbb{A}$, then $\mathbb{B}_\phi$ retracts on to the copy of $\mathbb{A}$ and so is a YES instance of CSP($\mathbb{A}$). Conversely, if $\mathbb{B}_\phi$ is a YES instance of CSP($\mathbb{A}$) (so there is a homomorphism $\psi : \mathbb{B}_\phi \to \mathbb{A}$), then as $\mathbb{B}_\phi$ contains $\mathbb{A}$ as an induced substructure, and $\mathbb{A}$ is a core, we can follow ψ by some automorphism α of $\mathbb{A}$ to ensure that $\alpha \circ \psi$ is a homomorphism that acts as the identity map on the substructure $\mathbb{A}$ of $\mathbb{B}_\phi$. Thus $\alpha \circ \psi$ extends ϕ, as required. Note that oracle queries can be assumed as automatically cleaned after each query, regardless of space.

Deciding separation and identification is similar. If we want to witness R-separation, then for each $(b_1, \ldots, b_k) \notin R^{\mathbb{B}}$ (of arity k) and each $(a_1, \ldots, a_k) \in S^{\mathbb{A}}$, we will perform the following check. If there is $b_j = b_{j'}$ with $a_j \neq a_{j'}$ then replace $(a_1, \ldots, a_k)$ with the next tuple in $\mathbb{A}$. Otherwise, we adjoin a copy of $\mathbb{A}$ to $\mathbb{B}$ and identify $b_1 := a_1$, $b_2 := a_2$, $\ldots$, $b_k := a_k$. This can be performed in logspace and written to an oracle for CSP($\mathbb{A}$). For each $(b_1, \ldots, b_k) \notin R_i^{\mathbb{B}}$ we need to discover one such tuple $(a_1, \ldots, a_k) \in S_i^{\mathbb{A}}$ for which $\mathbb{B}'$ is a YES instance of CSP($\mathbb{A}$). If there are none, reject. If so, then move to the next tuple $(b_1, \ldots, b_k) \notin R_i^{\mathbb{B}}$, each time wiping the tape clean.

Essentially the same technique works to verify R-identification: instead of systematically matching each $(b_1, \ldots, b_k) \notin R^{\mathbb{B}}$ (of arity k) with each $(a_1, \ldots, a_k) \in S^{\mathbb{A}}$, we match $(b_1, \ldots, b_k) \notin R^{\mathbb{B}}$ (of arity k) with $(a_1, \ldots, a_k) \in R^{\mathbb{A}}$. □

Example 3.3. *Recall that the complete graph $\mathbb{K}_2$ is the usual template for G2C (identically, monotone 1-in-2 2SAT). As $\mathbb{K}_2$ is a core structure and* CSP($\mathbb{K}_2$) *is decidable in logspace, then membership of finite structures in* $\mathsf{SP}(\mathbb{K}_2)$ *is also decidable in logspace, as is k-robust* 2*-colourability (for any fixed k). Similarly, if $\mathbf{A}$ denotes the template for directed graph unreachability, then* $\mathsf{SP}(\mathbf{A})$ *has membership problem decidable in nondeterministic logspace and k-robust unreachability is decidable in nondeterministic logspace.*

The assumption of being a core in Proposition 3.2 is necessary, at least when it comes to testing the separation conditions. In [19] a finite semigroup is given with NP-complete SP-membership. The semigroup can be built over any graph H and is denoted by T_H. For NP-completeness it suffices to use first graph in Example 2.4;

we here denote this graph by $\mathbb{K}_3^\sharp$. Let $\mathrm{graph}(T_{\mathbb{K}_3^\sharp})$ be the structure on the universe of $T_{\mathbb{K}_3^\sharp}$ with a single ternary relation: the graph of the binary operation of $T_{\mathbb{K}_3^\sharp}$.

Example 3.4. $\mathrm{CSP}(\mathrm{graph}(T_{\mathbb{K}_3^\sharp}))$ *is trivial, but the membership problem for* $\mathsf{SP}(\mathrm{graph}(T_{\mathbb{K}_3^\sharp}))$ *is* NP*-complete.*

Proof. We use $\mathbb{G}$ to abbreviate $\mathrm{graph}(T_{\mathbb{K}_3^\sharp})$ in this proof. Note that $\mathbb{G}$ has a single relation, which is ternary. The triple $(0,0,0)$ is a tuple in this relation. So any instance of $\mathrm{CSP}(\mathbb{G})$ may be mapped to 0 and have the fundamental relation preserved. Now we reduce G3C to $\mathsf{SP}(\mathbb{G})$. Given any graph H, let S_H denote the graph of the semigroup T_H. It is easy to verify that homomorphisms from S_H to $\mathbb{G} = S_{\mathbb{K}_3^\sharp} = \mathrm{graph}(T_{\mathbb{K}_3^\sharp})$ coincide with semigroup homomorphisms from T_H to $T_{\mathbb{K}_3^\sharp}$. However any homomorphism from T_H into $T_{\mathbb{K}_3^\sharp}$ separating 0 from f induces a 3-colouring of H (this fact is Lemma 5.1 of [19]). Then the fact that H is 3-colourable if and only if $T_H \in \mathsf{SP}(T_{\mathbb{K}_3^\sharp})$ gives $S_H \in \mathsf{SP}(\mathbb{G})$ as required. □

It is clear from the proof that membership in the class $\mathsf{RP}(\mathrm{graph}(T_{\mathbb{K}_3^\sharp}))$ is also NP-complete. An alternative example may be taken by taking the disjoint union of a single looped vertex with any simple graph $\mathbb{G}$ with known NP-complete finite membership problem for its universal Horn class (such as one described in part (1) of Example 2.4); we omit the details.

Part 2. Flexible satisfaction

4. Robust NAE3SAT

The main result in this section is Lemma 4.3, which is a technical lemma showing something stronger than the NP-completeness of 2-robust not-all-equal 3SAT.

Gottlob [14] showed that for any number $k > 0$, there is a reduction from 3SAT to $3(k+1)$-satisfiability, with the property that NO instances map to NO instances, but YES instances map to k-robustly satisfiable YES instances. In [2] it is shown that the standard reduction from SAT to 3SAT translates 3-robust 12SAT to the 2-robust 3SAT. The notion of robust satisfiability extends to graph 3-colourability in a natural way, and it is shown that 2-robust graph 3-colourability is NP-complete. This is proved using a result in Beacham [4] which (in the guise of "unfrozen 3-hypergraph 2-colourability") showed the NP-completeness of 2-robust 3-hypergraph 2-colourability.

There is a straightforward reduction from G3C to monotone 1-in-3 3SAT, however this reduction does not in general translate 2-robust 3-colourability to 2-robust monotone 1-in-3 3SAT. We give a fresh proof of the NP-completeness of 2-robust G3C, but where the constructed G3C are triangulated, in the sense that each edge lies within a triangle. This extra property is potentially of some interest in itself (the requirement arises in [19] for example), but in this article is used to facilitate an easy subsequent reduction to 2-robust monotone 1-in-3 3SAT. We mention that one can always add triangles to an existing graph and preserve 3-colourability, however such adjustments do not in general preserve 2-robust 3-colourability. We prove slightly more: whenever x, y, z are variables such that $(x \vee y \vee z)$ (or some permutation of this) does *not* appear as a clause in the instance, then we will find that at least one pair from x, y, z fails to appear in any clause of the instance, thus it is locally compatible to assign the value $(1, 1)$ to this pair, and therefore there

is a 1-in-3 satisfying truth assignment failing the missing clause $(x \vee y \vee z)$ in the 1-in-3 sense. Also, as the pair $(0,1)$ is always a locally compatible assignment for a variable pair (x,y) (with $x \neq y$), we find that it is NP-complete to recognise edge and equality unfrozen instances of monotone 1-in-3 3SAT. Equivalently, deciding membership in the universal Horn class of the template for monotone 1-in-3 3SAT has NP-complete membership problem for finite structures.

We begin by recalling the NP-completeness of k-robust $(3k+3)$-SAT and some of the details of Gottlob's reduction to this from 3SAT. Let I be an instance of 3SAT over the set of variables $X = \{x_1, \ldots, x_n\}$. For any k, consider the set $X_k = \{x_{i,j} \mid i \in \{1,\ldots,n\} \text{ and } j \in \{1,\ldots,2k+1\}\}$. Gottlob [14] constructs an instance of $(3k+3)$-SAT as follows: each clause $(\bar{x} \vee \bar{y} \vee \bar{z})$ in I (where $\bar{x}, \bar{y}, \bar{z}$ are elements of $\{x_1, \ldots, x_n, \neg x_1, \ldots, \neg x_n\}$) is replaced by the family of all $\binom{2k+1}{k+1}^3$ clauses of the form

$$(\bar{x}_{i_1} \vee \cdots \vee \bar{x}_{i_{k+1}} \vee \bar{y}_{i'_1} \vee \cdots \vee \bar{y}_{i'_{k+1}} \vee \bar{z}_{i''_1} \vee \cdots \vee \bar{z}_{i''_{k+1}})$$

where $\{i_1, \ldots, i_{k+1}\}$, $\{i'_1, \ldots, i'_{k+1}\}$, $\{i''_1, \ldots, i''_{k_{k+1}}\}$ are $(k+1)$-element subsets of $\{1, \ldots, 2k+1\}$. Here if $\bar{x} = x_i$ then by $\bar{x}_{i_j}$ we mean x_{i,i_j} and similarly for $\bar{y}$ and $\bar{z}$. And if $\bar{x} = \neg x_i$ then by $\bar{x}_{i_j}$ we mean $\neg x_{i,i_j}$, and so on. The corresponding instance of $(3k+3)$-SAT will be denoted by $I^\sharp$. It is easy to see that this reduction can be performed in linear time (albeit with a fairly large constant) and in logspace.

Each satisfying truth assignment ν for I extends to a satisfying truth assignment $\nu^\sharp$ for $I^\sharp$ by giving each $x_{i,j}$ the value of x_i. But in fact, it suffices to give the majority of the $x_{i,1}, \ldots, x_{i,2k+1}$ the value of x_i, because each $(k+1)$-subset of $\{1, \ldots, 2k+1\}$ will also include one of the elements in this majority. Conversely, each truth assignment ν for $I^\sharp$ gives rise to a truth assignment $\nu^\flat$ for I as follows: give $x_i \in X$ the majority value in $\nu(x_{i,1}), \ldots, \nu(x_{i,2k+1})$. (Yes, $(\nu^\sharp)^\flat = \nu$.) We have already seen that if $\nu^\flat$ satisfies I then ν satisfies $I^\sharp$. Now assume that $\nu^\flat$ fails to satisfy I. So some clause $(\bar{x} \vee \bar{y} \vee \bar{z})$ takes the value 0 under $\nu^\flat$. Let $\{i_1, \ldots, i_{k+1}\}$, $\{i'_1, \ldots, i'_{k+1}\}$, $\{i'_1, \ldots, i'_{k+1}\}$ be $(k+1)$-element subsets of $\{1, \ldots, 2k+1\}$ such that $\nu(\bar{x}_{i_1}) = \cdots = \nu(\bar{x}_{i_{(k+1)}}) = \nu^\flat(\bar{x})$ and $\nu(\bar{y}_{i'_1}) = \cdots = \nu(\bar{y}_{i'_{(k+1)}}) = \nu^\flat(\bar{y})$ and $\nu(\bar{z}_{i_1}) = \cdots = \nu(\bar{z}_{i_{(k+1)}}) = \nu^\flat(\bar{z})$, which exist due to the definition of $\nu^\flat$. Then the clause

$$(\bar{x}_{i_1} \vee \cdots \vee \bar{x}_{i_{k+1}} \vee \bar{y}_{i'_1} \vee \cdots \vee \bar{y}_{i'_{k+1}} \vee \bar{z}_{i''_1} \vee \cdots \vee \bar{z}_{i''_{k+1}})$$

exists in $I^\sharp$ and has the same truth value under ν as $(\bar{x} \vee \bar{y} \vee \bar{z})$ under $\nu^\flat$. Thus I is satisfiable if and only if $I^\sharp$ is satisfiable. But moreover if I is satisfiable, then $I^\sharp$ is k-robustly satisfiable. To see this, fix any satisfying truth assignment for I and let p be any partial assignment from some k-set of variables $Y \subseteq X_k$ to $\{0,1\}$. As $k < k+1$, we can adjust $\nu^\sharp$ to agree with p on all variables in Y without changing the majority value of $\nu^\sharp$. Thus p extends to a satisfying truth assignment for $I^\sharp$.

We now wish to observe that this proof continues to hold if we look for NAE-satisfaction. In particular, if we begin from monotone NAE3SAT, then we arrive at k-robust monotone NAE$(3k+3)$SAT. Thus we have the following result (which is a trivial, though useful, variation of Gottlob's argument in [14]).

Theorem 4.1. *For any $k > 0$, there is a logspace reduction from SAT to monotone NAE$(3k+3)$SAT such that no instances of SAT give rise to no instances of monotone NAE$(3k+3)$SAT, but yes instances of SAT give rise to k-robustly satisfiable yes instances of monotone NAE$(3k+3)$SAT.*

We mention that an alternative approach is to take k-robust $(3k+3)$SAT and add one further literal f to all clauses. This also yields the NP-completeness of k-robust NAE$(3k+4)$SAT, however it is not monotone. There is a technical step later in this article where it is useful for us to have a monotone instance of NAESAT (the value $(3k+4)$ or $(3k+3)$ is unimportant to us).

It is shown in [2] that the standard reduction from SAT to 3SAT translates the 3-robust 12SAT problem to the 3-robust 3SAT problem. We want something similar now for NAE3SAT. In fact we need only slightly more than 2-robustness, but additionally need extra information about which triples of literals can take which patterns of variables. Of particular interest to us is placing literals in majority values under some NAE-satisfying truth assignment.

The standard reduction from mSAT (for $m > 3$) to 3SAT replaces clauses $(\ell_1 \vee \ell_2 \vee \cdots \vee \ell_m)$ by the conjunction of 3-clauses

$$(\ell_1 \vee \ell_2 \vee z_2) \wedge (\neg z_2 \vee \ell_3 \vee z_3) \wedge \cdots \wedge (\neg z_{m-3} \vee \ell_{m-2} \vee z_{m-2}) \wedge (\neg z_{m-2} \vee \ell_{m-1} \vee \ell_m)$$

by way of the introduction of $m-3$ new variables $z_2, \ldots z_{m-2}$. Starting from $I^\sharp$ constructed above, we let the new instance of NAE3SAT be denoted by $f(I)$. Again this reduction can be computed in logspace. We refer to literals built from the introduced variables as the *Z-literals*, while those in X_k are referred to as *X-variables*. Note that the z_i are chosen afresh for each clause (so that a more careful though cumbersome nomenclature would include the clause name). We refer to $(\ell_1 \vee \cdots \vee \ell_m)$ as the *ancestor clause* for the displayed family of 3-clauses, and that the displayed family of 3-clauses are the descendants. We note that there is no loss in generality in assuming that the alphabet X_k of $I^\sharp$ is given some linear order, and that each clause of $I^\sharp$ lists its literals in agreement with this order. This also gives the descendant clauses a natural left-to-right direction (starting from small variables in X_k moving rightwards to increasing size). We do not order the Z-literals.

It is well known that this is a valid reduction from NAESAT to NAE3SAT. Nevertheless we give some details on how to translate between NAE-satisfying truth assignments for $I^\sharp$ and NAE-satisfying truth assignments for $f(I)$, as we will make detailed reference to these later.

Let $c = (y_1 \vee \cdots \vee y_m)$ be a clause of $I^\sharp$, and for each $i = 2, \ldots, m-2$, let z_i be the variable introduced in the $(i-1)^{\text{th}}$ descendant clause of c in the reduction of $I^\sharp$ to $f(I)$. For a given assignment on X-variables NAE-satisfying I, we may always extend to a truth assignment on the Z-variables so that $I^\sharp$ is NAE-satisfied. In this extension, it is possible that the truth values for some Z-variables can be chosen with complete freedom, however other variables may be forced to take a particular value. In order to prove 2-robust satisfiability of $I^\sharp$, we may be given a partial assignment p that gives a value to a Z-variable z, and want to call on the k-robust NAE-satisfiability of I to achieve an extension of p to $I^\sharp$. We need to ensure that the NAE-satisfying truth assignment we use does not require z to take a different value than p has given it. For this we initially extend p to some X-variables neighbouring z: but these in turn may force particular values for other Z-literals. To avoid such problems, the trick we employ is to extend p to a block of X-variables surrounding z in such a way that the Z-literals at the furthest extremities of this block have flexible choice. This turns out to be easy. For the Z variable z_i (introduced for clause c), we use by default the X-variables $y_{i-2}, y_{i-1}, y_i, y_{i+1}, y_{i+2}$, however on occasions we need to take the union of two such sets when we are considering two

distinct variables of the form of z_i: these variables may also appear in the same family of descendant clauses, and we may need to expand our range of consideration slightly further to allow a consistent choice of values for locally appearing variables. Fortunately, there is lots of flexibility.

Lemma 4.2. *Let $(y_1 \vee \cdots \vee y_m)$ be an ancestor clause in I. Fix $2 \leq i \leq m-2$ and let $j, j' \geq 0$ be such that $i - j \geq 1$, $i + j' < m$ and not both $i - j = 1$ and $i + j' = m$. Let $p : \{z_i, y_{i-j}, \ldots, y_{i+j'}\} \to \{0,1\}$ with $p(y_{i-j}) \neq p(y_{i-j+1})$ and $p(y_{i+j'-1}) \neq p(y_{i+j'})$. Then for any value of $p(z_{i-j}), p(z_{i+j'}) \in \{0,1\}$ it is possible to extend p to the variables $z_{i-j+1}, \ldots, z_{i+j'-1}$ without violating NAE-satisfaction. Note that if $i - j = 1$ then z_{i-j} does not exist and statement refers only to $z_{i+j'}$, and dually if $i + j' = m - 1$ then $z_{i+j'}$ does not exist, and the statement refers only to z_{i-j}.*

Proof. This follows quite easily: it is always easy to extend truth values to the z-variables in the descendant clauses, provided that when the first or the final clause in the family is reached, if the two y-variables have the same value, then the inherited truth value of the z variable is not the same. The assumptions in the lemma ensure this has not happened. Moreover, when there is a change in values to $p(y_{i-j+1}) \neq p(y_{i-j})$ there either emerges a choice of truth value at z_{i-j+1} or at z_{i-j}. In either case there is a choice at z_{i-j}. A dual statement holds for $z_{i+j'}$. □

We mention that for a single Z-variable z_i we only ever need select the values of y_i and y_{i+1}: choose the value of y_i differently from the value given z_i, and choose y_{i+1} to be given the same truth value as z_i (hence different to $\neg z_i$). For the instance $\cdots \wedge (\ell \vee y_i \vee z_i) \wedge (\neg z_i \vee y_{i+1} \vee \ell') \wedge \cdots$:

$$\text{for } z_i = 0: \qquad \cdots \wedge (\ell \vee 1 \vee 0) \wedge (1 \vee 0 \vee \ell') \wedge \cdots \tag{3}$$

$$\text{for } z_i = 1: \qquad \cdots \wedge (\ell \vee 0 \vee 1) \wedge (0 \vee 1 \vee \ell') \wedge \cdots \tag{4}$$

The displayed clauses are NAE-satisfied with free choice remaining for the values of ℓ and ℓ'. However sometimes we will already have inherited a value for one or both of y_i and y_{i+1}, and these may conflict. The worst case is where both have pre-determined values, and they are opposites to what is described in items (3) or (4). If $i > 2$ so that z_{i-1} exists, then this forces the value of z_{i-1} (to agree with z_i), and similarly if $i < m-1$ so that z_{i+1} exists, then the value of z_{i+1} is also forced. However, in this case we can simply move to y_{i-1} and y_{i+2} in a similar way. For the instance

$$\cdots (\ell \vee y_{i-1} \vee z_{i-1}) \wedge (\neg z_{i-1} \vee y_i \vee z_i) \wedge (\neg z_i \vee y_{i+1} \vee z_{i+1}) \wedge (\neg z_{i+1} \vee y_{i+2} \vee \ell') \cdots$$

we have:

$$\text{for } z_i = 0: \qquad \cdots (\ell \vee 1 \vee 0) \wedge (1 \vee 0 \vee 0) \wedge (1 \vee 1 \vee 0) \wedge (1 \vee 0 \vee \ell') \cdots \tag{5}$$

$$\text{for } z_i = 1: \qquad \cdots (\ell \vee 0 \vee 1) \wedge (0 \vee 1 \vee 1) \wedge (0 \vee 0 \vee 1) \wedge (0 \vee 1 \vee \ell') \cdots \tag{6}$$

Again the displayed clauses are NAE-satisfied with free choice remaining for the values of ℓ and ℓ'. The situation where $i = 2$ or $m-1$ is very similar but easier, as one fewer Z-literal is involved.

Typically applications of Lemma 4.2 will use $j = j' = 2$ and assign the variables $y_{i-2}, y_{i-1}, y_i, y_{i+1}, y_{i+2}$ truth values alternating between 0 and 1 (but any pattern that has a different truth value for y_{i-2}, y_{i-1} and a different truth value for y_{i+1}, y_{i+2} works equally well). When $i = 2$, then apply alternating values to y_1, y_2, y_3, y_4

(though an alternative is to choose the value of y_1 and y_2 to be the same, but the opposite to the any desired value for z_2), and dually at the right hand end.

We often will refer to this process as "stabilising" the selected value of a Z literal: making a selection of neighbouring X-variable truth values consistent with the given value of the Z-literal, and finishing with a flexible choice of the truth value of the Z-literals appearing at the boundary of the interval of descendant clauses.

Lemma 4.3. *The mapping $I^\sharp \mapsto f(I)$ from monotone NAE(21)SAT to NAE(3)SAT is a logspace reduction with the additional property that if $I^\sharp$ is 6-robustly NAE-satisfiable, then $f(I)$ has the following properties*:

1. *$f(I)$ is 2-robustly NAE-satisfiable.*
2. *If ℓ and ℓ' are literals appearing in distinct clauses c and d of $f(I)$, then any locally compatible partial truth assignment $p : \{\ell, \ell'\} \to \{0, 1\}$ extends to a NAE-satisfying truth assignment for $f(I)$ such that ℓ takes a majority value in c. Note that ℓ may coincide with ℓ' or be its negation.*
3. *if z is a Z-literal, with z in clause c and $\neg z$ in clause d, then there are X-variables x, y in c and d respectively such that $f(I)$ can be NAE-satisfied by an assignment p with $p(x) = p(z)$ and $p(y) = p(\neg z)$.*
4. *for each clause $(\ell \vee \ell' \vee \ell'')$ appearing in $f(I)$, every partial assignment NAE-satisfying the clause $(\ell \vee \ell' \vee \ell'')$ extends to a full NAE-satisfying truth assignment.*

Proof. Let p be a locally compatible partial truth assignment on 2 literals ℓ, ℓ' of the constructed NAE3SAT instance.

If both ℓ, ℓ' are X-variables, then we require only 2-robustness of $I^\sharp$ to enable a NAE-satisfying truth assignment of $I^\sharp$ extending p, which then gives a corresponding NAE-satisfying truth assignment of $f(I)$ that also extends p. (Obviously this can also be done for up to $n = 6$ of the X-variables, a fact we make use of below.)

Now assume that ℓ is an X-variable (say, x) but ℓ' is a Z-literal. Apply Lemma 4.2, assigning alternating values to the X-variables immediately surrounding z. If this does not conflict with the given value of x, then we can use the 6-robust satisfiability of $I^\sharp$ to extend to a NAE-satisfying truth assignment. If it does, then use the dual alternating alternating assignment (which will not conflict with the value of x).

So finally (for item (1)), assume that both ℓ and ℓ' are Z-literals. First select the two X-variables appearing in a clause with ℓ (if ℓ appears in a boundary clause, then select one X-variable in the boundary clause, and select the only X-variable in the other clause containing ℓ or $\neg\ell$). Let us denote these by x, x'. First use (3) or (4) to give values to x, x'. Now we must do the same for ℓ': as a first attempt we can try (3), (4), however as observed above, it is possible that the current assignment for x, x' conflicts with that being asked by (3), (4). In this case it is easy to move to the next X-variables across, as in (5) and (6). There is one minor variation to this: where ℓ' is a Z-literal in a clause also containing ℓ or $\neg\ell$. This is also easily covered using a very basic variation of (3). In each case though, we need extend p to at most four X-variables. Then p can be extended to a NAE-satisfying truth assignment on $I^\sharp$, then to a NAE-satisfying truth assignment for $f(I)$ consistent with the given values $p(\ell)$ and $p(\ell')$. Thus we have 2-robust NAE-satisfaction (item 1).

Now to show (2) and (3). Part (3) is a strengthening of (2) for one particular case, so we really just prove (2), but with the extra information proved in the case to which (3) applies. Let $p : \{\ell, \ell'\} \to \{0, 1\}$ be arbitrary.

If $\ell = \ell'$ then because of the assumption that ℓ and ℓ' appear in different clauses, it follows that ℓ is an X-variable, and moreover that c and d do not arise from the same ancestor clause of $I^\sharp$. Select a Z-literal z appearing in c. Set $p(z) := p(\ell)$. In order to stabilise this choice, we need to select to further X-variables: an X-variable x in the unique clause containing z' and an X-variable y in the unique clause containing $\neg z$. By setting $p(x) := p(\neg z')$ and $p(y) := p(z)$, we can use 3-robust satisfiability of $I^\sharp$ to extend the values of p on ℓ, x, y to a full NAE-satisfying assignment for $I^\sharp$, whose extension to a NAE-satisfying truth assignment for $f(I)$ agrees with the values p gives z, z'.

If $\ell = \neg\ell'$ then we have the conditions of (3) and we must find both ℓ and ℓ' in majority configuration under some NAE-satisfying truth assignment. Only Z-variables appear in negated form, and Z-variables appear in exactly two clauses (one in positive form, one in negative form). Let x be an X-variable in the one clause c containing ℓ and y an X-variable in the one clause d containing $\neg\ell$. We give x the same value as ℓ and y the same value as $\neg\ell$. This forces the value of the remaining literal in c as well as the remaining literal in d. It is easy to see following Lemma 4.2 that we can assign values for at most one further X variable (each) on either side of x and y to stabilise this forcing and enable flexible assignment to further Z-variables. (Of course, if one of c or d is a boundary clause this extra choice is not required.) This gives (3) (and hence (2) for this case).

Now assume that ℓ and ℓ' are related to different variables.

If both ℓ and ℓ' are X-variables, then we can select a Z-literal z in the clause c. If $\neg z$ occurs in clause d, then replace the choice of z with the other literal in c (this could be an X-variable or a Z-literal, it doesn't matter: if it is an X-variable it cannot be ℓ'). Thus we may assume that $\neg z$ does not appear in clause d. Give z the same value as ℓ, and assign values to some neighbouring X-variables to stabilise the choice of value for z, accommodating the value for ℓ' if it appears here. Using this construction, at most four X-variables will be required: as well as ℓ and ℓ', we need two variables to stabilise the chosen value of z. We mention that even when ℓ' appears in a neighbouring clause to ℓ, there are still only two further X-variables that need to be given fixed values (note that z' and z'' are Z-variables and x, y are X-variables):

$$\cdots \wedge (\ldots \vee x \vee z'') \wedge (\neg z'' \vee \ell' \vee z') \wedge (\neg z' \vee \ell \vee z) \wedge (\neg z \vee x \vee \ldots) \wedge \cdots$$

If, for example, $p(\ell) = p(\ell') = 0$, then the instruction is to set $p(z) = 0$, which forces $p(z') = 0$, which in turn forces $p(z'') = 0$. This is stabilised by adding $p(x) = 1$ and $p(y) = 0$. If however $p(\ell) = 0 \neq p(\ell')$ then only three X-variables are required: the selection of $p(z) = 0$ still forces $p(z') = 0$, but this stabilises in the clause containing ℓ', so that x is not required. The variable y is still required.

Now assume ℓ is an X-variable, but ℓ' is not. The technique is essentially identical: select a literal z in c that does not appear in positive or negative form in d; we will set $p(z) = p(\ell)$, but we must show that this extends to a full NAE-satisfying assignment. If z can be chosen to be an X-variable (meaning c was a boundary clause) then this option should be taken. If this is possible, then we need only select two extra X-variables to stabilise the given value of ℓ'. However if z is a Z-literal, we must be slightly more careful. We want can stabilise the choice of $p(z) = p(\ell)$ by adding at most two variables, however selecting values for these variables may affect our ability to stabilise the given value $p(\ell')$. If d is not a boundary clause,

then we may simply select two further X-variables around d to stabilise the selection of any Z-literals. However, if d is a boundary clause, then it is possible that the variables selected to stabilise the given value of z both appear in clause d and cause it to fail. In this instance, we can just replace z with the other Z-literal in c (given that we arrive at this situation only when c was not a boundary clause). Thus in every case, at most two extra variables are required to stabilise the value given for $p(\ell')$. In total there are up to five X-variables required here: ℓ along with two to stabilise the value for $p(z)$ and two to stabilise the value for $p(\ell')$.

Next assume that ℓ is a Z-literal and ℓ' is an X-variable. This is similar, but easier. Select an X-variable x in clause c. If x is not ℓ', or if $x = \ell'$ but $p(\ell') = p(\ell)$, then we can set $p(x) = p(\ell)$ and only two further X-variables need be selected to stabilise the given value of $p(\ell)$. If x is ℓ' and $p(\ell) \neq p(\ell')$, we can select either a different X-variable in d (if d is a boundary clause) or the other Z-literal, z' say. Two further X-variables are required to stabilise this case, making a total of four X-variables in total.

Finally, assume that both ℓ and ℓ' are Z-variables. In this case, we simply choose neighbouring values around them, ensuring that the truth value of the selected X-variable in clause c agrees with the value of ℓ. To stabilise both the given value of $p(\ell)$ and $p(\ell')$, a further four X-variables need to be assigned values, thought fewer are required in some cases. Four new X-variables are required, for example, when x appears in clause d and $p(x) = p(\ell) = p(\ell')$.

This completes the proof of (2) (while (3) was incorporated into a case of (2)). Now we prove (4). This is quite self evident though: simply select one or two X variables either side of such a clause to stabilise the value of any Z-literals: at most 2 extra X-variables need be assigned. □

5. Edge and vertex robust G3C

The main result of this section is the following theorem, the first statement of which is Theorem 3.3 of [2]. In Proposition 5.2 we also observe a slight extension to this theorem, that shows in the case that our construction is 3-colourable, then there is also no non-local restrictions to the number of colours required to cover two distinct edges.

Theorem 5.1. *2-robust graph 3-colourability is NP-complete. Moreover, this remains true for graphs in which every edge lies within a triangle.*

Proof. The construction is closely based on a standard reduction of NAE3SAT to G3C except we have more triangles, and the robust colourability is proved by way of the technical conditions in Lemma 4.3.

We are given an instance $f(I)$ as constructed in Section 4 and must construct a graph $\mathbb{G}$. The vertices are as follows.

- A1 There is a special vertex a.
- B1 For each variable x there are vertices corresponding to both x and $\neg x$ (the *literal vertices*).
- C1 For each clause e we have three vertices e_1, e_2, e_3 (the *clause position vertices*).
- D1 For each clause e we have three connector vertices $c_{e,1}, c_{e,2}, c_{e,3}$.

The edges are described as follows.

- E1 The special vertex is adjacent to all literal vertices.

F1 Each literal vertex x is adjacent to its negation $\neg x$.
G1 For each clause e, the clause vertices e_1, e_2, e_3 form a clique.
H1 For each clause $e = (\ell_1 \vee \ell_2 \vee \ell_3)$ and each $i \in \{1, 2, 3\}$, the following three vertices form a 3-clique: the literal vertex z_i, the connector vertex $c_{e,i}$ and ith clause position vertex e_i for e.

Observe that the property that no clause contains both a variable and its negation ensures that there are no 4-cycles: more precisely there are no four distinct vertices u_0, u_1, u_2, u_3 with each of $\{u_i, u_{i\oplus_4 1}\}$ an edge.

We first recall that the constructed graph $\mathbb{G}$ is 3-colourable if and only if the NAE3SAT instance I is satisfiable in the NAE sense. Moreover, whenever there is a 3-colouring of $\mathbb{G}$ in which the special vertex is coloured 2, then the colours of the literal vertices lie amongst $\{0, 1\}$ and necessarily correspond to a NAE-satisfying truth assignment. We omit details of this claim as this is completely standard and also routinely verified. Thus $\mathbb{G}$ is 3-colourable if and only if the NAE3SAT instance I is NAE-satisfiable. Now assume that I is NAE-satisfiable, and hence is 3-robustly NAE-satisfiable and has the other properties described in Lemma 4.3.

We first verify 2-robustness of $\mathbb{G}$. By symmetry it suffices to consider only 3-colourings where the special vertex is coloured by 2. Let u, v be a pair of distinct but nonadjacent vertices. It suffices to like-colour u and v and to unlike-colour u and v (then choosing colour names achieves all possible colourings of u, v).

Note that each vertex aside from the special vertex is associated with a unique literal: a literal vertex is associated with itself, a connector vertex is associated with the unique literal it is adjacent to, and similarly a clause position vertex with the literal it is adjacent to. For a vertex w we let $\lceil w \rceil$ denote this literal. Let $\lfloor w \rfloor$ denote the clause position vertex adjacent or equal to w, and let $[w]$ denote the connector vertex adjacent or equal to w. Similarly, each vertex aside from the special vertex and the literal vertices is associated with a unique clause. For such a vertex w we let e_w denote the clause associated with the vertex w.

Case 1. One of u, v is the special vertex (say, $u = a$).
As v is not adjacent to u, it follows that u is either a connector vertex or a clause position vertex. By the 2-robustness of I, we may select a partial truth assignment giving $\lceil u \rceil$ a majority value $j \in \{0, 1\}$ in the clause e_u. Then we can colour $\lfloor u \rfloor$ by 2 or by $1 - j$, which in turn colours the connector vertex $[u] = c_{e,\lceil u \rceil}$ by $1 - j$ or 2, respectively. These two colourings can be used to like- and unlike-colour u and v. This completes Case 1.

Case 2. u and v are literal vertices.
As u and v are not adjacent, they correspond to different variables x, y in I. By the 2-robustness of I we may find truth values giving these the same or different values. This corresponds to the same or different colours in $\mathbb{G}$. This completes Case 2.

Case 3. u is a literal vertex and v is a connector vertex or a clause position vertex.
Let $x = \lceil u \rceil$ (which actually gives $u = x$ here) and $y = \lceil v \rceil$. As u and v are not adjacent, we cannot have $x = y$, but we could have $x = \neg y$.

Subcase 3.1. If $x = \neg y$. Let z be any literal with an occurrence in e_v, aside from y. The 2-robustness of I enables us to construct a NAE-satisfying truth assignment in which y and z take the same truth value 0. In this case x is assigned the value 1, so that u is coloured 1. The colour of the vertex $\lfloor v \rfloor$ is a choice of either colour 2

or the colour 1 (the opposite to the truth value of y). These colourings will colour $[v]$ by either 1 or 2, respectively. Thus we may like- or dislike-colour u and v.

Subcase 3.2. x is not adjacent to y.

Subsubcase 3.2.1. x appears in the clause e_v. So the clause corresponding to e_v is $(x \vee y \vee z)$ for some literal z. Then by item (4) of Lemma 4.3, there is a NAE-satisfying truth assignment giving (x, y, z) the value $(1, 0, 0)$. The corresponding 3-colourings have a choice of $\lfloor v \rfloor$ taking either the same colour as u or the colour 2, and the same (but in reverse order) for $[v]$.

Subsubcase 3.2.2. $\neg x$ appears in the clause e_v. So e_v is the clause $(\neg x \vee y \vee z)$ and item (4) of Lemma 4.3 shows that there is a NAE-satisfying truth assignment for $f(I)$ that gives the value $(0, 0, 1)$ to $(\neg x, y, z)$. The corresponding 3-colourings give $\lfloor v \rfloor$ either the same colour 0 as u or the colour 2, and the same (but in reverse order) for $[v]$.

Subsubcase 3.2.3. Now assume that neither x nor $\neg x$ appears in $e_v = (z \vee y \vee z')$ (for some literals z, z'). Find a truth value giving x the opposite truth value as y, and with y taking the majority value in the clause e_v; this is possible by Lemma 4.3 part (2). This produces 3-colourings for which u is coloured by the truth value of x, while v is either coloured 2 or the negation of the truth value of y (with flexible choice of either case). In other words, we have a choice to colour u and v differently or the same (respectively).

Case 4. u and v are both connector vertices.

A connector vertex will be coloured 2 unless they connect to a clause position vertex whose position was in majority value, in which case there is a choice of 2 or the negation of the truth value of the literal in that position. For two connector vertices u and v, we may make these choices independently (thus like- or dislike-colour freely) except in the case where of where u and v connect to the same clause component. In this situation we give $\lceil u \rceil$ and $\lceil v \rceil$ different truth values, and observe that one of $\lceil u \rceil$ and $\lceil v \rceil$ (say, $\lceil u \rceil$) will take a majority value in the corresponding clause e_u. The corresponding graph colouring will have v coloured 2, but there is a choice of either colouring u by 2 or by the negation of the truth value of $\lceil u \rceil$. Thus we may like colour or dislike colour u and v.

Case 5. u is a connector vertex and v is a clause position vertex.

Note that as u and v are not adjacent, we do not have $\lfloor u \rfloor = v$.

Case 5.1. $\lceil u \rceil = \lceil v \rceil$.

In this case, e_u and e_v are different clauses. Select a literal z in e_u but not equal to $\lceil u \rceil$, nor to $\neg\lceil u \rceil$. Using 2-robustness, consider a truth assignment giving $\lceil u \rceil$ and z the same truth value. The corresponding colouring gives a choice of colours for u: either the negation of the truth value of z, or 2. Similarly v is coloured by the negation of the truth value of z, or possibly a choice of colour 2. It is clear that u and v may be like- and dislike-coloured.

Case 5.2. $\lceil u \rceil = \neg\lceil v \rceil$.

In this case, e_u and e_v are different clauses. By Lemma 4.3 part (3), there is NAE-satisfying truth assignment giving both $\lceil u \rceil$ and $\lceil v \rceil$ a majority value in the clauses e_u and e_v. In the corresponding 3-colouring we can choose to give both $\lfloor u \rfloor$ and $\lfloor v \rfloor = v$ the colour 2 (which unlike-colours u from v) or to give u the colour 2 and v the colour 2, which like-colours the vertices.

Case 5.3. $\lceil u \rceil$ and $\lceil v \rceil$ are literals built over distinct variables.

First assume that $\lfloor u \rfloor$ and $\lfloor v \rfloor$ are in different clause components. Consider a truth

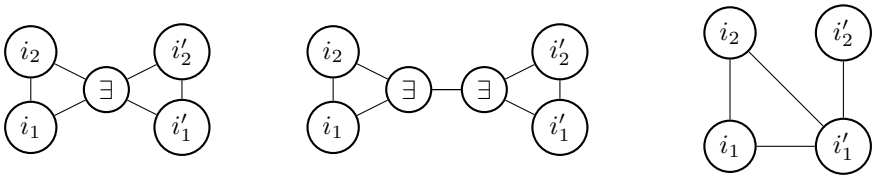

FIGURE 2. Three local configurations forcing the number of colours on $\{i_1, i_2, i_1', i_2'\}$ in any valid 3-colouring (other edges may also be present). In the first, exactly two colours can appear in $\{i_1, i_2, i_1', i_2'\}$. In the second and third, exactly 3 colours appear in $\{i_1, i_2, i_1', i_2'\}$. In the third configuration, we allow i_2' to coincide with one of i_1 or i_2.

assignment giving $\lceil u \rceil$ and $\lceil v \rceil$ the same value, but also such that $\lceil v \rceil$ has the majority truth value in e_v (this is possible by Lemma 4.3 part (2)). Then as we can choose give the colour 2 to the connector vertex u, and the negation of the truth value of $\lceil v \rceil$ to v, we can unlike-colour. But we can also choose to give colour 2 to v because $\lceil v \rceil$ has the majority truth value in e_v.

Now assume that $e_u = e_v$. In this case we can easily like- or unlike-colour u and v by a suitable truth assignment identifying or distinguishing $\lceil u \rceil$ and $\lceil v \rceil$. Identifying the value of $\lceil u \rceil$ and $\lceil v \rceil$ forces u and v to have the same colour (both 2, or both the negation of the truth value of $\lceil u \rceil$). Distinguishing their truth values gives a choice of colouring to one of the vertices, according to whether $\lceil u \rceil$ or $\lceil v \rceil$ is in majority configuration (and one must be).

Case 6. u and v are clause position vertices.
As u and v are not adjacent we must have $e_u \neq e_v$. If we can give $\lceil u \rceil$ and $\lceil v \rceil$ the same truth value, then it is easy to see that we can like-colour u and v. If however $\lceil u \rceil = \neg \lceil v \rceil$, then we instead use Lemma 4.3 part (3) to obtain a NAE-satisfying truth assignment in which $\lceil u \rceil$ takes a majority value in e_u and $\lceil v \rceil$ takes a majority value in e_v. Then we can like-colour u and v by the colour 2, as well as dislike colour, by letting u and v take the alternative choice (the negation of their truth values, which are distinct given the assumption $\lceil u \rceil = \neg \lceil v \rceil$).

So it remains to dislike-colour u and v in the case where $\lceil u \rceil \neq \neg \lceil v \rceil$. If $\lceil u \rceil = \lceil v \rceil$, then select some other literal z in e_v and give it the same truth value as $\lceil v \rceil$, which can be done using the 2-robust NAE-satisfiability of I. Then we can colour v by 2, keeping the colour of u fixed as the negation of the truth value of $\lceil v \rceil$ (thus we can dislike-colour u and v). Otherwise, $\lceil u \rceil$ and $\lceil v \rceil$ are built over distinct variables. Select a truth assignment giving $\lceil u \rceil$ and $\lceil v \rceil$ different truth values. Then we obtain a colouring where u and v have different colours (possibly both can be chosen to have colour 2, but this choice is not forced: we can select them to be coloured by the negation of the truth value of the corresponding literal, different for u and v). This completes Case 6, the final case.

This completes the proof of the 2-robust 3-colourability of $\mathbb{G}$. □

We can observe more. Let $\{i_1, i_2\}$ and $\{i_1', i_2'\}$ be two edges, not identical. There are always at least two colours appearing amongst $\{i_1, i_2, i_1', i_2'\}$, but there are some simple local configurations that when present can force there to be always exactly 2 colours, or always exactly 3 colours.

Proposition 5.2. *Let $\{i_1, i_2\}$ and $\{i'_1, i'_2\}$ be two edges, not identical. Except in the three configurations given in Figure 2 (and permutations of these labellings), if there is a 3-colouring, then there is a 3-colouring giving $\{i_1, i_2, i'_1, i'_2\}$ exactly 2 colours and a 3-colouring giving $\{i_1, i_2, i'_1, i'_2\}$ exactly 3 colours.*

Proof. We show that in all other cases we may achieve both the situation of exactly two colours amongst $\{i_1, i_2, i'_1, i'_2\}$ and three colours amongst $\{i_1, i_2, i'_1, i'_2\}$.

Case 1. When $|\{i_1, i_2, i'_1, i'_2\}| = 3$ but $\{i_1, i_2, i'_1, i'_2\}$ is not a 3-clique: say, $i_1 = i'_1$, but $\{i_2, i'_2\}$ is not an edge. In this case we can use 2-robust 3-colourability to like-colour i_1 and i_2 to achieve exactly two colours in the set, and dislike-colour i_1 and i_2 to achieve all three colours.

Case 2. $|\{i_1, i_2, i'_1, i'_2\}| = 4$ and but $\{i_1, i_2, i'_1, i'_2\}$ does not contain a 3-clique, nor forms a pair of 3-cliques with some fifth vertex. In this case there are at most two edges between $\{i_1, i_2\}$ and $\{i'_1, i'_2\}$.

Subcase 2.1. There are two edges between $\{i_1, i_2\}$ and $\{i'_1, i'_2\}$.
As there is no 3-clique, there is no loss of generality in assuming that the edges are $\{i_1, i'_1\}$ and $\{i_2, i'_2\}$. But this is a 4-cycle, and there are no such configurations in $\mathbb{G}$. So there is nothing to check.

Subcase 2.2. There is exactly one edge between $\{i_1, i_2\}$ and $\{i'_1, i'_2\}$; say $\{i_1, i'_1\}$. By like colouring i_2 and i'_2 we have all three colours amongst $\{i_1, i_2, i'_1, i'_2\}$. So the goal is to achieve exactly 2 colours amongst the four vertices. Let i_3 and i'_3 be vertices forming a 3-clique with $\{i_1, i_2\}$ and $\{i'_1, i'_2\}$ respectively. It is clear from inspection of $\mathbb{G}$ that i_3 cannot coincide with i'_3 (alternatively observe that this case has already been considered as one where 3 colours *must* be used). Also i_3 is not amongst $\{i'_1, i'_2\}$, nor is i'_3 amongst $\{i_1, i_2\}$, because there is only one edge between the vertices of the two edges being considered in Subcase 2.2. If i_3 and i'_3 are *not adjacent*, then by like-colouring them we achieve exactly two colours amongst our four vertices. However, if i_3 and i'_3 are adjacent then $i_3 \sim i'_3 \sim i'_1 \sim i_1 \sim i_3$ is a four cycle, which does not appear in $\mathbb{G}$. Thus Subcase 2.2 is complete.

Subcase 2.3. There are no edges between $\{i_1, i_2\}$ and $\{i'_1, i'_2\}$.
Again let i_3 and i'_3 be vertices forming 3-cliques with $\{i_1, i_2\}$ and $\{i'_1, i'_2\}$ respectively. As there are no edges between $\{i_1, i_2\}$ and $\{i'_1, i'_2\}$, it follows that $i_3 \notin \{i'_1, i'_2\}$ and $i'_3 \notin \{i_1, i_2\}$. Also $i_3 \neq i'_3$ as this is a configuration where we already know exactly two colours can be found amongst $\{i_1, i_2, i'_1, i'_2\}$ under any valid 3-colouring. Now if $\{i_3, i'_3\}$ is an edge then we have already observed that exactly 3 colours appear amongst $\{i_1, i_2, i'_1, i'_2\}$. Thus we may assume that $\{i_3, i'_3\}$ is not an edge. But then we can colour i_3 and i'_3 identically to place exactly two colours appear amongst $\{i_1, i_2, i'_1, i'_2\}$, and colour i_3 and i'_3 differently to place exactly 3 colours amongst $\{i_1, i_2, i'_1, i'_2\}$ as required. □

5.1. **Some comments.** To finish this section, we make a brief comment about the prospect of following the direct approach to 2-robust G3C sketched in [2]. In the next section we require every edge to lie in a triangle (technically this is not essential: we can add new variables instead, but these variables correspond in a natural way to the vertices we would otherwise have added, and the issues we now describe continue to arise for these variables). It is easy to adjust the construction in [2] (which is a standard one) to add triangles to any untriangulated edges, in essentially the same way as we have done in the present section. Then we also obtain analogues of the connector vertices. However under colourings giving the special vertex 2, connector vertices always take the value 2 unless they connect to a clause position

taking a majority value (for [2] this is a 3-hyperedge, but these are just monotone 3-clauses). A potential problem now arises in like-colouring a variable vertex v of the hypergraph (which encodes as itself in the G3C instance) with a connector vertex. As v only ever takes the colour 0 and 1 (assuming, as always that the special vertex is given the colour 2), we require that the connector vertex to be connecting to a clause position taking a majority value, and that this majority value is the opposite of the truth value of the variable v. This is slightly stronger than 2-robustness: it requires an extension of a partial assignment on two variables, but with the additional property that a specific one of them lies in a majority configuration in some given clause. This small issue is the cause of the technicalities in Lemma 4.3. The author would be very surprised if something similar to Lemma 4.3 failed to hold true for Beacham's construction, however verifying this did not appear to be a shorter path to the present point of this article: Beacham's construction is at least as complicated as the one we present above and additionally builds off the top of another rather technical reduction.

The author's original approach to 2-robust 3-colourability was by a reduction from the Hamiltonian Circuit problem (via an intermediate encoding into SAT then NAE3SAT). The idea is that if there is a Hamiltonian circuit in a graph, then we can also choose where to start this circuit, and in which direction to travel. By varying these choices it is possible to again arrive at a reduction to G3C, in which yes instances are 2-robustly satisfiable. However, the author was unable to push this to the results of the next section (due to essentially the same problem as just identified above) until a conversation with Standa Živný at the ALC Fest in Prague 2014 revealed the more flexible approach of Gottlob [14], which is where the present approach begins.

6. 2-robust monotone 1-in-3 3SAT and separability.

We now encode the graph $\mathbb{G}$ into monotone 1-in-3 3SAT; the main result is Theorem 6.1, which was the primary target of all previous sections in Part 2. Recall from Section 1 that $\mathbb{2}$ denotes the usual template for monotone 1-in-3 3SAT as a constraint problem: the domain is the set $\{0,1\}$ and there is a single ternary relation $R = \{(0,0,1),(0,1,0),(1,0,0)\}$.

The variables are straightforward: for each vertex i of $\mathbb{G}$ and each colour $j = 0,1,2$, the variable $x_{i,j}$ (which asserts "vertex i is coloured j").

We now give the clauses.

- A2 for each vertex i we have the clause $(x_{i,0} \vee x_{i,1} \vee x_{i,2})$ (and the other 5 permutations of this clause). 1-in-3 satisfaction of this clause corresponds to "vertex i is coloured by exactly one colour".
- B2 for each 3-clique $\{i_1,i_2,i_3\}$ and each colour $j = 0,1,2$, we have the clause $(x_{i_1,j} \vee x_{i_2,j} \vee x_{i_3,j})$ (and the other 5 permutations of this clause). 1-in-3 satisfaction of this clause corresponds to "colour j appears exactly once in the 3-clique $\{i_1,i_2,i_3\}$".

Because every edge lies within a 3-clique, it is easy to see that 1-in-3 satisfying truth assignments of these clauses coincide precisely to valid 3-colourings of $\mathbb{G}$ under the given interpretation of the variables. We now extablish 2-robust satisfaction of monotone 1-in-3 3SAT in the case that $\mathbb{G}$ is 3-colourable (using the fact that for the particular graphs constructed, this additionally implies 2-robust 3-colourability).

Let x and y be two distinct variables; say $x = x_{i,j}$ and $y = x_{i',j'}$. We must show that every locally compatible assignment of values to x and y extends to a 1-in-3 satisfying truth assignment.

Case 1: $i = i'$. In this case, local consistency ensures that not both x and y can take the value 1 (due to clauses of kind A2). Without loss of generality, assume that $j = 0$ and $j' = 1$. In this case, as $x \neq y$ we can easily range through the values $(0,0), (0,1), (1,0)$ for (x,y) by colouring vertex i by 2, 1 and 0 respectively.

Case 2: $i \neq i'$. Note that distinct vertices of $\mathbb{G}$ can always be dislike coloured. Thus, if $j \neq j'$ (say, $j = 0$, $j' = 1$) then we may achieve the following values of (x,y) in the following ways:

- $(0,0)$ by colouring i by 1 and i' by 0.
- $(0,1)$ by colouring i by 2 and i' by 1.
- $(1,0)$ by colouring i by 0 and i' by 2.
- $(1,1)$ by colouring i by 0 and i' by 1.

If $j = j'$ (say, $j = 0$), then the possibilities depend on whether or not $\{i, i'\}$ is an edge. If it is, then local consistency with clauses of kind B2 ensures we do not need to achieve (x,y) taking the value $(1,1)$. We may achieve the following values of (x,y) in the following ways:

- $(0,0)$ by colouring i by 1 and i' by 2.
- $(0,1)$ by colouring i by 1 and i' by 0.
- $(1,0)$ by colouring i by 0 and i' by 1.
- $(1,1)$ by colouring i by 0 and i' by 0, which is possible when $\{i, i'\}$ is *not* an edge. When $\{i, i'\}$ is an edge, then assigning both x and y the value 1 is not locally compatible with the corresponding clause of kind B2.

Theorem 6.1. *Let 2 denote the usual template for monotone 1-in-3 3SAT. Any set of monotone 3SAT instances containing the 2-robustly monotone 1-in-3 satisfiable members of* $\mathsf{SP}(2)$ *and contained within the members of* $\mathrm{CSP}(2)$ *has* NP*-hard membership with respect to logspace many-one reductions.*

Proof. The chain of reductions in the article so far maps (by a logspace many-one reduction) YES instances of the NP-complete problem monotone NAE3SAT to the 2-robustly monotone 1-in-3 satisfiable members of $\mathsf{SP}(2)$ and NO instances to NO instances of $\mathrm{CSP}(2)$. □

In the next section we need a number of particular properties held by the instances of monotone 1-in-3 3SAT that are obtained via this reduction. Before we list them, let us say that two pairs of variables $\{v, w\}$ and $\{x, y\}$ are *linked* if there is a variable z such that $(v \vee w \vee z)$ and $(x \vee y \vee z)$ are clauses appearing in I.

(I) If $(x \vee y \vee z)$ is a clause that does *not* appear amongst the constructed clauses, then if the instance is satisfiable, there is a satisfying truth assignment giving $(x \vee y \vee z)$ a value with at least two occurrences of truth value 1. (There is no requirement that $|\{x, y, z\}| = 3$ here: when, say $x = y$ this is easy as we simply find a truth assignment making x true.)

(II) No two distinct clauses share more than one variable. This does not require the instance to be a YES instance.

(III) Consider the graph whose vertex set is the set of all two element sets of variables appearing in I, and where the edge relation is between linked pairs. Then this graph is a disjoint union of cliques. This does not require the instance to be satisfiable.

(IV) If w, x, y, z are four variables (not necessarily distinct), and the instance is a YES instance, then there is a 1-in-3 satisfying truth assignment ν such that at least two elements of the sequence $\nu(w), \nu(x), \nu(y), \nu(z)$ are 1.

For Observation (I), note that if $(x \vee y \vee z)$ fails to be in the construction, then for at least one of the pairs (x, y), (x, z), (y, z) it is locally compatible to assign the value $(1, 1)$. This follows easily by inspection of the clause construction and 2-robust 1-in-3 satisfiability.

Observation (II) follows immediately from the construction of the clauses.

Observation (III) follows from Observation (II) and the fact that all clauses in the instance have exactly 3 variables. Note that every pair of variables x, y that appear in some clause, the pair $\{x, y\}$ is in a nontrivial component of the graph described in Observation (III).

Observation (IV), will follow immediately from Observation (I) unless both $(w \vee x \vee y)$ and $(x \vee y \vee z)$ are clauses in the instance. But then Observation (II) shows that these clauses must be identical (up to a permutation of the position of variables in the clause), so that $w = z$. Then any 1-in-3 satisfying truth assignment with $\nu(z) = 1$ yields $\nu(w), \nu(x), \nu(y), \nu(z)$ equal to $1, 0, 0, 1$.

Part 3. Applications to variety membership

7. Variety membership for $\mathbf{B}_2^1$

Recall that the variety generated by an algebraic structure $\mathbf{A}$ is the class of all algebras of the same signature that satisfy the equations holding on $\mathbf{A}$; equivalently that are homomorphic images of subalgebras of direct powers of $\mathbf{A}$. The pseudovariety of $\mathbf{A}$ is the class of finite algebras in the variety of $\mathbf{A}$, but more generally a pseudovariety is a class of finite algebras of the same signature that is closed under taking homomorphisms, subalgebras and finitary direct products. The finite part of any variety is a pseudovariety, but not every pseudovariety is of this form. The computational problem of deciding membership of finite algebras in a variety coincides with the problem of deciding membership in the pseudovariety of finite members. For this reason, the results of the present section are cast in the setting of pseudovarieties.

The six element monoid $\mathbf{B}_2^1 = \langle 1, a, b \mid a^2 = b^2 = 0, aba = a, bab = b \rangle$ is the most ubiquitous minimal counterexample in problems relation to varieties of semigroups. The possible intractability of the computational problem of deciding membership of finite semigroups in the variety generated by $\mathbf{B}_2^1$ is perhaps the most obvious unresolved problem relating to $\mathbf{B}_2^1$ and so it is not surprising that this has appeared in a number of places in the literature including Problem 4 of Almeida [1, p. 441], Problem 3.11 of Kharlampovich and Sapir [23] and page 849 of Volkov, Gol′dberg and Kublanovksiĭ [38]. We now use the results of the previous section to show that this computational problem is NP-hard. In fact we show this true for a large swathe of finite semigroups: any semigroup whose pseudovariety contains that of $\mathbf{B}_2^1$ and lies within the join of the pseudovariety of $\mathbf{B}_2^1$ with the pseudovariety $\mathbb{L}\mathbf{DS}$.

We build on the reduction to 2-robust monotone 1-in-3 3SAT in the previous section, though we now reuse the symbol I to denote a typical instance created from this reduction. The notation V_I denotes the set of variables appearing in I. Recall that if $(x \vee y \vee z)$ is a clause in I, then so are the other five permutations of

this clause. Also, no variable appears twice in any clause, and no pair of distinct clauses contain two common variables.

For each instance I we construct a semigroup S_I. The semigroup S_I is generated from the following elements:

(1) 1, a multiplicative identity element;
(2) 0, a multiplicative 0 element;
(3) an element $\mathbf{a}$ and an element $\mathbf{b}$;
(4) for each variable v of I an element $\mathbf{a}_v$.

The semigroup is subject to the following rules:

(1) $\mathbf{aba} = \mathbf{a}$, $\mathbf{bab} = \mathbf{b}$, $\mathbf{aa} = \mathbf{bb} = 0$ (so that $1, \mathbf{a}, \mathbf{b}$ generate a copy of $\mathbf{B}_2^1$);
(2) $\mathbf{a}_u\mathbf{a}_v = 0$ if there is no clause containing both u and v (or if $u = v$).
(3) $\mathbf{a}_u\mathbf{a}_v = \mathbf{a}_v\mathbf{a}_u$.
(4) $\mathbf{a}_u\mathbf{a}_v\mathbf{a}_w = \mathbf{a}$ whenever $u \vee v \vee w$ is a clause in I.
(5) $\mathbf{a}_{u_1}\mathbf{a}_{u_2} = \mathbf{a}_{v_1}\mathbf{a}_{v_2}$ if $\{u_1, u_2\}$ and $\{v_1, v_2\}$ are linked pairs of variables.

When u and v appear in some common clause, then we let $[uv]$ denote the set of all pairs xy where x and y are variables and $\{u, v\}$ is linked to $\{x, y\}$; so $[uv] = [xy]$ in this case. Of course, $[xy] = [yx]$ simply because $\{x, y\} = \{y, x\}$ (provided x and y appear together in some clause of I). We will write $\mathbf{a}_{[uv]}$ to denote the value $\mathbf{a}_u\mathbf{a}_v$ (and note that $\mathbf{a}_{[uv]} = \mathbf{a}_{[vu]}$ by condition (3)). Note that if $\{u, v\}$ is linked to $\{x, y\}$, then $[uv] = [xy]$, which is consistent with the property $\mathbf{a}_{[uv]} = \mathbf{a}_u\mathbf{a}_v = \mathbf{a}_x\mathbf{a}_y = \mathbf{a}_{[xy]}$ implied by condition (5).

If $u \vee v \vee w$ is a clause in I then so are the further 5 permutations of the variables, which is consistent with the commutativity condition (3). Also, for any variable u we have $\mathbf{a}_u\mathbf{a}_u = 0$ by (2), because of the property that the variable u does not appear twice in any single clause. Finally, we observe that for any quadruple of variables v_1, v_2, v_3, v_4 we have $\mathbf{a}_{v_1}\mathbf{a}_{v_2}\mathbf{a}_{v_3}\mathbf{a}_{v_4} = 0$. Indeed, if $(v_1 \vee v_2 \vee v_3)$ is not a clause, then it is observed already in Section 6 that one of the pairs $\{v_1, v_2\}, \{v_1, v_3\}, \{v_2, v_3\}$ does not appear in any clause, thus the product will be 0 by condition (2) and the commutativity condition (3). Thus to avoid $\mathbf{a}_{v_1}\mathbf{a}_{v_2}\mathbf{a}_{v_3}\mathbf{a}_{v_4} = 0$ we require $(v_1 \vee v_2 \vee v_3)$ to be a clause in I. The same argument shows that $(v_2 \vee v_3 \vee v_4)$ must be a clause in I. But no two distinct clauses share two common variables. Thus these clauses coincide and hence $|\{v_1, v_2, v_3, v_4\}| = 3$ and conditions (2) and (3) again yield 0 as the value of the product.

Let $\mathbf{c} := \mathbf{ab}$ and $\mathbf{d} := \mathbf{ba}$. Thus the elements of S_I other than $0, 1$ are the following:

- $\mathbf{b}$, $\mathbf{a}$, $\mathbf{c}$, $\mathbf{d}$,
- $\mathbf{a}_u$, $\mathbf{ba}_u$, $\mathbf{ca}_u$, $\mathbf{a}_u\mathbf{b}$, $\mathbf{a}_u\mathbf{d}$, $\mathbf{a}_u\mathbf{ba}_v$ for all variables u, v,
- $\mathbf{a}_{[uv]}$, $\mathbf{ba}_{[uv]}$, $\mathbf{ca}_{[uv]}$, $\mathbf{a}_{[uv]}\mathbf{b}$, $\mathbf{a}_{[uv]}\mathbf{d}$, $\mathbf{a}_{[uv]}\mathbf{ba}_w$, $\mathbf{a}_w\mathbf{ba}_{[uv]}$ for all variables u, v, w for which both u and v appear in a clause of I.
- $\mathbf{a}_{[u_1u_2]}\mathbf{ba}_{[v_1v_2]}$ for all variables u_1, u_2, v_1, v_2 for which u_1, u_2 both appear in some single clause and v_1, v_2 both appear in some single clause.

We mention that the equality in condition (5) of the presentation does not lead to further equalities between the remaining elements just listed: this is due to Observation (III) (which is independent of the satisfiability of I) and the equality (4) of the presentation. In particular the semigroup S_I can be constructed in logspace from the instance I, given that the instance I satisfies the particular properties identified above.

Lemma 7.1. *If I is robustly 1-in-3 satisfiable, then $S_I \in \mathbb{V}(\mathbf{B}_2^1)$.*

Proof. Let T be the subalgebra of $(\mathbf{B}_2^1)^{\hom(I,2)}$ generated by the following maps

$$\beta, \alpha, \alpha_v, 1\colon \hom(I,2) \to \mathbf{B}_2^1$$

for each $v \in V_I$:

$$\begin{aligned} \beta(\phi) &= b \\ \alpha(\phi) &= a \\ \alpha_v(\phi) &= \begin{cases} 1 & \text{if } \phi(v) = 0 \\ a & \text{if } \phi(v) = 1 \end{cases} \\ 1(\phi) &= 1. \end{aligned}$$

Let $U \subseteq T$ consist of all elements t of T that contain a 0-entry; that is, there is $\phi \in \hom(I,2)$ with $t(\phi) = 0$. The set U is trivially seen to be an ideal of T, and we claim that the Rees quotient T/U is isomorphic to S_I under the map ι defined by $\alpha_v \mapsto \mathbf{a}_v$, $\beta \mapsto \mathbf{b}$ and $1 \mapsto 1$. (Note that this will give $\alpha \mapsto \mathbf{a}$ because of rule (4) of the definition of S_I.) To prove this we first verify that ι is well defined on the generators: this only requires that there are no equalities between the elements 1, β, α and the various α_v. This is trivially true between 1, β and α. For each variable v in I, the 2-robust satisfiability of I ensures there are homomorphisms $\nu_1, \nu_2 : I \to 2$ with $\nu_1(v) = 1$ and $\nu_2(v) = 0$. Then $\alpha_v(\nu_1) = a$ and $\alpha(\nu_2) = 1$. So $\alpha_v \notin \{1, \alpha, \beta\}$. Similarly, for distinct variables u, v it is locally compatible to satisfy I by a homomorphism ν with $\nu(u) \neq \nu(v)$, and then $\alpha_u(\nu) \neq \alpha_v(\nu)$. This shows that ι maps the generators of T/U bijectively to those of S_I.

We now verify that the various relations defining S_I hold for the corresponding elements of T/U; it will then remain to show that elements distinct in S_I correspond to distinct elements of T/U.

The relations (1) and (3) are trivial to verify. For relations of kind (2), observe that when u, v are variables not both appearing in the same clause, then the partial assignment $(u, v) \mapsto (1, 1)$ is locally compatible, so there is $\nu \in \hom(I,2)$ with $\nu(u) = \nu(v) = 1$. Then on coordinate ν we have both $\alpha_u(\nu) = a$ and $\alpha_v(\nu) = a$ so that $[\alpha_u\alpha_v](\nu) = 0$, showing that $\alpha_u\alpha_v \in U$ (or is 0 in T/U). For relations of kind (4), note that the definition of 1-in-3 satisfaction ensures that exactly one variable in each clause $(x \vee y \vee z)$ takes the value 1, which in turn implies that exactly one of $\alpha_x(\nu), \alpha_y(\nu), \alpha_z(\nu)$ takes the value a, with the others taking the value 1. Thus $[\alpha_x\alpha_y\alpha_z](\nu) = a$ (for all $\nu \in \hom(I,2)$), showing that $\alpha_x\alpha_y\alpha_z = \alpha$ when $(x \vee y \vee z)$ is a clause in I.

For the rules in item (5), consider $\{u_1, u_2\}$ and $\{v_1, v_2\}$ such that $(w \vee u_1 \vee u_2)$ and $(w \vee v_1 \vee v_2)$ are clauses in I. Let $\nu \in \hom(I,2)$ be arbitrary. If $\nu(w) = 1$, then $\nu(u_1) = \nu(u_2) = \nu(v_1) = \nu(v_2) = 0$ and $\alpha_{u_1}\alpha_{u_2}(\nu) = \alpha_{v_1}\alpha_{v_2}(\nu) = 1$. If $\nu(w) = 0$, then $\{\nu(u_1), \nu(u_2)\} = \{\nu(v_1), \nu(v_2)\} = \{0, 1\}$ and $\alpha_{u_1}\alpha_{u_2}(\nu) = \alpha_{v_1}\alpha_{v_2}(\nu) = a$. As $\nu \in \hom(I,2)$ is arbitrary, we have $\alpha_{u_1}\alpha_{u_2} = \alpha_{v_1}\alpha_{v_2}$. As for elements of S_I, we now write $\alpha_{[uv]}$ to denote the product $\alpha_u\alpha_v$.

Next we show that distinct elements in S_I correspond to distinct elements in T/U. This has already been performed for the generators. Observe that each element α_u has a coordinate equal to 1. Similarly, if u, v appear in some common clause of I, then, as $(u, v) \mapsto (0, 0)$ is locally compatible, there is a coordinate

$\nu \in \hom(I,(2))$ with $\alpha_u\alpha_v(\nu) = 1$. Thus elements of the form α_u or $\alpha_u\alpha_v$ (where u, v appear in some clause of I) are distinct from any product involving α, β or $\alpha_x\alpha_y\alpha_z$ where $x \vee y \vee z$ is a clause.

We have already verified that if $u \neq v$ then $\alpha_u \neq \alpha_v$. Observe now that if u and v are distinct variables appearing in some clause of I and w is a variable, then without loss of generality we have $w \neq u$, so that $(w, u) \mapsto (0, 1)$ is locally compatible. Hence there is a coordinate $\nu \in \hom(I,(2))$ in which $\alpha_u\alpha_v(\nu) = a \neq 1 = \alpha_w(\nu)$. Hence no α_w is equal to any $\alpha_u\alpha_v$.

Next, if u, v appear in some clause and x, y appear in some clause but $\{u, v\}$ and $\{x, y\}$ are not linked. Let w be such that $(w \vee u \vee v)$ is a clause in I. Now $(w \vee x \vee y)$ is not a clause, so by Observation (I) there is $\nu \in \hom(I, 2)$ failing to 1-in-3 satisfy $(w \vee x \vee y)$. Because x and y do appear in some clause of I it follows that at most one of $\nu(x)$ and $\nu(y)$ is 1. Because $(w \vee x \vee y)$ fails under ν we must have $(\nu(w), \nu(x), \nu(y)) \in \{(1, 0, 1), (1, 1, 0)\}$. Then $\nu(u) = \nu(v) = 0$ so that $\alpha_{[uv]}(\nu) = 1$ while $\alpha_{[xy]}(\nu) = a$. Thus $\alpha_{[uv]} \neq \alpha_{[xy]}$ if $[uv] \neq [xy]$.

It is now easy to verify that there are no further equalities amongst the elements of T/U aside than for corresponding elements of S_I. For example, because $\alpha_u \neq \alpha_v$ when $u \neq v$, any nonzero product $\alpha_u\beta\ldots$ will be distinct from $\alpha_v\beta\ldots$, because there is a coordinate ν at which $\alpha_u(\nu) = a$ and $\alpha_v(\nu) = 1$, and then $[\alpha_u\beta\ldots](\nu) \in \{a, ab\}$ while $[\alpha_v\beta\ldots](\nu) \in \{b, ba\}$. While there are many such cases, they are all trivial.

This shows that the map ι is a homomorphism from T/U onto S_I which is injective on $\iota^{-1}(S_I\backslash\{0\})$. However, Observation (IV) shows that 0 is the only element of T/U outside of $\iota^{-1}(S_I\backslash\{0\})$, so that ι is an isomorphism, as claimed. □

The converse to Lemma 7.1 is true and implies that deciding membership in the pseudovariety of $\mathbf{B}_2^1$ is NP-hard. We show something stronger than the converse. The class of all finite semigroups whose variety does *not* contain $\mathbf{B}_2^1$ turns out to form a pseudovariety: in fact it is the pseudovariety $\mathbb{L}\mathbf{DS}$ of finite semigroups which are locally in the pseudovariety $\mathbf{DS}$. It can be axiomatised by the pseudo-identity

$$(x^\omega y_1 x^\omega z x^\omega y_2 x^\omega)^\omega \approx ((x^\omega y_1 x^\omega z x^\omega y_2 x^\omega)^\omega z (x^\omega y_1 x^\omega z x^\omega y_2 x^\omega)^\omega)^\omega$$

(where, as usual, the notation x^ω denotes the limit $\lim_{n\to\infty} x^{n!}$) though we do not use this fact below. We use $\mathbb{V}(\mathbf{B}_2^1)$ to denote the pseudovariety generated by $\mathbf{B}_2^1$ (it consists of the finite members of the variety generated by $\mathbf{B}_2^1$). The interval $[\![\mathbb{V}(\mathbf{B}_2^1), \mathbb{L}\mathbf{DS} \vee \mathbb{V}(\mathbf{B}_2^1)]\!]$ consists of all pseudovarieties $\mathscr{V}$ with $\mathbb{V}(\mathbf{B}_2^1) \subseteq \mathscr{V} \subseteq \mathbb{L}\mathbf{DS} \vee \mathbb{V}(\mathbf{B}_2^1)$, or equivalently all pseudovarieties $\mathscr{V} \subseteq \mathbb{L}\mathbf{DS} \vee \mathbb{V}(\mathbf{B}_2^1)$ containing $\mathbf{B}_2^1$.

The proof of the next result makes use of Green's relations $\mathscr{H}$, $\mathscr{L}$, $\mathscr{R}$, $\mathscr{J}$. We briefly recall the definitions, but direct the reader to any text on semigroup theory for further details; Howie [17] is one such. Let S be a semigroup, and let S^1 denote the result of adjoining an identity element to S if it does not already have one (otherwise $S^1 = S$). Define the preorder $\leq_{\mathscr{L}}$ on S by $a \leq_{\mathscr{L}} b$ if there exists $x \in S^1$ with $xb = a$. The preorder $\leq_{\mathscr{R}}$ is defined dually, while $\leq_{\mathscr{H}}$ is the intersection $\leq_{\mathscr{L}} \cap \leq_{\mathscr{R}}$. The preorder $\mathscr{J}$ is defined by $a \leq_{\mathscr{J}} b$ if there are $x, y \in S^1$ with $xby = a$. The relations $\mathscr{L}$, $\mathscr{R}$, $\mathscr{H}$, $\mathscr{J}$ are the equivalence relations of elements equivalent under the respective preorder. Thus, for example $a \mathscr{J} b$ if $a \leq_{\mathscr{J}} b$ and $b \leq_{\mathscr{J}} a$. We observe that if $a \leq_{\mathscr{H}} b$ and $b \leq_{\mathscr{J}} a$ in a finite semigroup, then $a \mathscr{H} b$.

Lemma 7.2. *Let $\mathscr{V}$ be any pseudovariety in the interval $[\![\mathbb{V}(\mathbf{B}_2^1), \mathbb{LDS} \vee \mathbb{V}(\mathbf{B}_2^1)]\!]$. If $S_I \in \mathscr{V}$ then I is 1-in-3 satisfiable.*

Proof. The proof works for any class $\mathscr{V}$ with $\mathbb{V}(\mathbf{B}_2^1) \subseteq \mathscr{V} \subseteq \mathbb{LDS} \vee \mathbb{V}(\mathbf{B}_2^1)$ (not necessarily a pseudovariety). As $S_I \in \mathscr{V}$, we also have $S_I \in \mathbb{LDS} \vee \mathbb{V}(\mathbf{B}_2^1)$, so there are finite semigroups $\mathbf{S}_1, \ldots, \mathbf{S}_k$ from $\mathbb{LDS}$ such that there is $T \leq \mathbf{S}_1 \times \cdots \times \mathbf{S}_k \times (\mathbf{B}_2^1)^\ell$ and a surjective homomorphism $\psi : T \twoheadrightarrow S_I$. As $\mathbb{LDS} \vee \mathbb{V}(\mathbf{B}_2^1)$ is closed under finitary direct products, there is no loss of generality in assuming that $k \leq 1$. Also, clearly $\ell > 0$ because $S_I \notin \mathbb{LDS}$ due to the fact S_I embeds $\mathbf{B}_2^1$. For each variable v appearing in I, select any element $\hat{a}_v \in \psi^{-1}(\mathbf{a}_v)$. Select an arbitrary element $\hat{e}$ from $\psi^{-1}(1)$ (unless one wishes to perform this proof in the monoid signature, and then the canonical choice of $\hat{e}$ is the constant tuple equal to 1 on each projection). As S_I is generated by the $\mathbf{a}_v$ and $\mathbf{b}$ and 1, we can assume that T is generated by the elements $\hat{a}_v$ along with $\hat{e}$ and any chosen element from $\psi^{-1}(\mathbf{b})$. We make a more careful selection of an element $\hat{b}$ from $\psi^{-1}(\mathbf{b})$.

To select $\hat{b}$, choose any $b' \in \psi^{-1}(\mathbf{b})$ that is minimal amongst elements of $\psi^{-1}(\mathbf{b})$ with respect to the $\mathscr{J}$ order. Recall that for any clause $(u \vee v \vee w)$ in I we have $\mathbf{b}\mathbf{a}_u\mathbf{a}_v\mathbf{a}_w\mathbf{b} = \mathbf{b}$, so that $\hat{b}\hat{a}_u\hat{a}_v\hat{a}_w\hat{b} \in \psi^{-1}(\mathbf{b})$. Now $\hat{b}\hat{a}_u\hat{a}_v\hat{a}_w\hat{b} \leq_{\mathscr{H}} \hat{b}$, but also by the $\mathscr{J}$-minimality of the choice of $\hat{b}$ and the fact that $\hat{b}\hat{a}_u\hat{a}_v\hat{a}_w\hat{b} \in \psi^{-1}(\mathbf{b})$, we have $\hat{b}\hat{a}_u\hat{a}_v\hat{a}_w\hat{b} \geq_{\mathscr{J}} \hat{b}$. Thus $\hat{b}\hat{a}_u\hat{a}_v\hat{a}_w\hat{b} \mathscr{H} \hat{b}$. As this is true in T, it is also true of any projection π from T into $\mathbf{B}_2^1$. As $\mathbf{B}_2^1$ is $\mathscr{H}$-trivial, it follows that $\hat{b}\pi = \hat{b}\hat{a}_u\hat{a}_v\hat{a}_w\hat{b}\pi$ always. Select any element $\hat{e}$ from $\psi^{-1}(1)$. A trivial variation of the argument just given shows that $\hat{e}\hat{b}\hat{e}\pi = \hat{b}\pi$ for projections into $\mathbf{B}_2^1$.

We want to show that for some such projection π into $\mathbf{B}_2^1$, we have $\hat{b}\pi \in \{a, b\}$. Now, $1, \mathbf{b}, \mathbf{a}$ generate a subsemigroup of S_I that is isomorphic to $\mathbf{B}_2^1$. Let $u \vee v \vee w$ be an arbitrary clause in I and let $\hat{a}$ denote the value $\hat{a}_u\hat{a}_v\hat{a}_w$ (the value of $\hat{a}$ may depend on the clause selected). The homomorphism ψ maps the subsemigroup T' of T generated by $\hat{e}, \hat{b}, \hat{a}$ onto the subsemigroup of S_I isomorphic to $\mathbf{B}_2^1$ (on $\{1, \mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, 0\}$). Now $\mathbf{B}_2^1$ has a unique minimal set of generators: the elements $a, b, 1$. Also, all proper subsemigroups of $\mathbf{B}_2^1$ lie in $\mathbb{LDS}$. So for some projection π we have $\{\hat{e}\pi, \hat{b}\pi, \hat{a}\pi\} = \{a, b, 1\}$; otherwise T' lies in the pseudovariety generated by $\mathbf{S}_1$ and a family of proper subsemigroups of $\mathbf{B}_2^1$; in other words, T' lies in $\mathbb{LDS}$, contradicting the fact that $\mathbf{B}_2^1$ is isomorphic to a quotient of T'. Therefore at least one such projection π has $\{\hat{e}\pi, \hat{b}\pi, \hat{a}\pi\} = \{a, b, 1\}$. As $\hat{e}\hat{b}\hat{e}\pi = \hat{b}\pi$ it follows that $\hat{b}\pi \in \{a, b\}$ in this case, and up to an obvious automorphism of $\mathbf{B}_2^1$, there is no loss of generality in assuming that $\hat{b}\pi = b$.

Now, for every variable u, there are variables v, w such that $(u \vee v \vee w)$ is a clause in I, and then the property

$$\mathbf{b}\mathbf{a}_u\mathbf{a}_v\mathbf{a}_w\mathbf{b} = \mathbf{b}\mathbf{a}_u\mathbf{a}_w\mathbf{a}_v\mathbf{b} = \mathbf{b}\mathbf{a}_v\mathbf{a}_u\mathbf{a}_w\mathbf{b} = \mathbf{b}\mathbf{a}_v\mathbf{a}_w\mathbf{a}_u\mathbf{b} = \mathbf{b}\mathbf{a}_w\mathbf{a}_u\mathbf{a}_v\mathbf{b} = \mathbf{b}\mathbf{a}_w\mathbf{a}_v\mathbf{a}_u\mathbf{b} = \mathbf{b}$$

on projection π gives $\hat{a}_u\pi \in \{a, 1\}$. Define a truth assignment by $u \mapsto 1$ if $\hat{a}_u\pi = a$ and $u \mapsto 0$ if $\hat{a}_u\pi = 1$. Then for each clause $(u \vee v \vee w)$ in I, the property $[\hat{b}\hat{a}_u\hat{a}_v\hat{a}_w\hat{b}]\pi = b$ ensures that exactly one of $\hat{a}_u\pi, \hat{a}_v\pi, \hat{a}_w\pi$ is a, or equivalently, exactly one of u, v, w is 1. Thus the clauses are satisfied in the 1-in-3 sense, as required. □

Theorem 7.3. *Let $\mathcal{V}$ be any pseudovariety in the interval $[\![\mathbb{V}(\mathbf{B}_2^1), \mathbb{LDS} \vee \mathbb{V}(\mathbf{B}_2^1)]\!]$. Then membership in $\mathcal{V}$ is NP-hard with respect to logspace reductions.*

Proof. If I is a 2-robustly 1-in-3 satisfiable instance of monotone 1-in-3 3SAT, then as $S_I \in \mathbb{V}(\mathbf{B}_2^1)$ by Lemma 7.1, we also have $S_I \in \mathcal{V}$. Conversely, if $S_I \in \mathcal{V}$ then Lemma 7.2 shows that I is 1-in-3 satisfiable. The result now follows from Theorem 6.1. □

Let us say that a completely 0-simple semigroup is *block diagonal* if its sandwich matrix may be arranged in block diagonal form (with each rectangular block consisting of strictly nonzero entries). Every orthodox completely 0-simple semigroup is block diagonal, but there are non-orthodox block diagonal completely 0-simple semigroups. It is not hard to verify that block diagonal completely 0-simple semigroups are precisely those satisfying the law $xv \neq 0 \ \& \ uv \neq 0 \ \& \ uy \neq 0 \Rightarrow xy \neq 0$.

The monoid obtained by adjoining an identity element to a completely 0-simple semigroup will be called a *completely* 0*-simple monoid*. The following result and its proof do not depend on whether or not we ask that the identity element be given as a distinguished element or if we treat these monoids in the semigroup signature.

Corollary 7.4. *Let* $\mathbf{M}$ *be a block diagonal completely 0-simple semigroup. If* $\mathbf{M}$ *is not completely regular* (*equivalently, if it has nontrivial zero divisors*)*, then the membership problem for finite semigroups or monoids in* $\mathbb{V}(\mathbf{M}^1)$ *is* NP*-hard.*

Proof. First observe that the pseudovariety of $\mathbf{M}$ contains $\mathbf{B}_2^1$: it is not in the pseudovariety $\mathbb{L}\mathbf{DS}$ because it is not completely regular. Second, note that $\mathbf{M}$ lies within the join of $\mathbb{V}(\mathbf{B}_2^1)$ with the pseudovariety of all finite completely regular semigroups, which is a subclass of $\mathbb{L}\mathbf{DS}$. Indeed, it is not hard to show that for suitable $n > 1$ we have $\mathbf{M}$ is a divisor of the direct product of $\mathbf{B}_n$ with a completely 0-simple semigroup without zero divisors. Let P be a sandwich matrix for some representation of $\mathbf{M}$ as a Rees matrix semigroup. Then $\mathbf{M}^1$ is a divisor of the direct product of $\mathbf{B}_n^1$ with completely 0-simple monoid without zero divisors. Because $\mathbb{V}(\mathbf{B}_n^1) = \mathbb{V}(\mathbf{B}_2^1)$, this shows that $\mathbf{M}^1$ lies in the join of a subpseudovariety of $\mathbb{L}\mathbf{DS}$ with $\mathbf{V}(\mathbf{B}_2^1)$ so that Theorem 7.3 applies. □

The five element semigroup $\mathbf{A}_2$ has presentation $\langle a, b \mid aba = a, bab = b, aa = 0, bb = b\rangle$, and with an identity element adjoined is denoted $\mathbf{A}_2^1$. This is a completely 0-simple monoid lying outside of the block diagonal completely 0-simple monoid of Corollary 7.4; moreover it lies in the variety of any finite completely 0-simple monoid that is not block diagonal. In a subsequent article [22] it will be shown that $\mathcal{V}(\mathbf{B}_2^1)$ is not finitely based within the pseudovariety $\mathcal{V}(\mathbf{A}_2^1)$, and that Lemma 7.2 fails to extend to include $\mathcal{V}(\mathbf{A}_2^1)$. So the methods of the present article do not directly shed any light on the complexity of deciding membership in $\mathcal{V}(\mathbf{A}_2^1)$. Thus the following problem is of immediate interest.

Problem 7.5. *Can membership in the pseudovariety* $\mathcal{V}(\mathbf{A}_2^1)$ *be decided in polynomial time?*

We wish to make some observations relating to Problem 7.5. Recall that a finite algebra $\mathbf{A}$ of finite signature $\mathscr{L}$, the *equational complexity function* $\beta_{\mathbf{A}} : \mathbb{N} \to \mathbb{N}$ is defined by $\beta_{\mathbf{A}}(n) = 1$ if all $\mathscr{L}$-algebras of size n lie in the variety of $\mathbf{A}$ and otherwise $\beta_{\mathbf{A}}(n)$ is defined to be the maximum over all n-element $\mathscr{L}$-algebras $\mathbf{B}$ not in the variety of $\mathbf{A}$, of the smallest equation satisfied by $\mathbf{A}$ and failing on $\mathbf{B}$. (The precise definition of $\beta_{\mathbf{A}}$ depends on the precise definition of "size" of an equation, but any sensible will suffice for the following observation.) If $\mathbf{A}$ has a finite basis for its equations, then $\beta_{\mathbf{A}}$ is bounded, and the possible converse to this statement remains

an important unsolved problem which is known to be equivalent to a well-known open problem of Eilenberg and Schützenberger [9]; see [31] for a detailed discussion. To date there are very few algebras $\mathbf{A}$ for which $\beta_{\mathbf{A}}$ is known not to bounded by a polynomial. We wish to observe that $\mathbf{A}_2^1$ either has non-polynomial bounded equational complexity, or has co-NP-easy membership problem (with respect to many-one reductions), or possibly both. The idea is on page 247 of Jackson and McNulty [20]: if the equational complexity of $\mathbf{A}_2^1$ is bounded by a polynomial $p(n)$, then given a finite semigroup $\mathbf{B}$ not in the variety of $\mathbf{A}_2^1$, we may guess an equation $u \approx v$ of length at most $p(|B|)$ that is true on $\mathbf{A}_2^1$, as well as an assignment of its variables into $\mathbf{B}$ witnessing its failure on $\mathbf{B}$. This certificate for non-membership in the variety of $\mathbf{A}_2^1$ may be verified in deterministic polynomial time using [24, Lemma 8] to verify that $u \approx v$ holds on $\mathbf{A}_2^1$ (as well as the easy check that the witnessing assignment into $\mathbf{B}$ does yield failure of $u \approx v$ on $\mathbf{B}$). We believe this provides reasonable evidence that membership in the pseudovariety of $\mathbf{A}_2^1$ is not NP-hard with respect to polynomial time many-one reductions.

Following Lee and Zhang [29] it is known that there are precisely 4 semigroups of order 6 that have no finite basis for their equations. Aside from $\mathbf{B}_2^1$, which we have shown to generate a pseudovariety with hard membership problem, and $\mathbf{A}_2^1$ which we have discussed in Problem 7.5, there is the example $\mathbf{A}_2^g$ observed by Volkov (private communication; see [38]) and the semigroup $\mathbf{L}$ of Zhang and Luo [40]. The semigroup $\mathbf{A}_2^g$ was shown to have polynomial time membership problem for its pseudovariety by Goldberg, Kublanovsky and Volkov [38], but the complexity of membership in $\mathcal{V}(\mathbf{L})$ is unknown.

Problem 7.6. *Can membership in $\mathcal{V}(\mathbf{L})$ be decided in polynomial time?*

At around the same time as the author proved the above results for $\mathbf{B}_2^1$, Klíma, Kunc and Polák distributed a manuscript that gave a family of semigroups for which membership problem in the variety generated by any one is co-NP-complete (with respect to many-one reductions) [25]. Moreover, like $\mathbf{B}_2^1$, these semigroups are known important examples in the theory of finite semigroups, generating the pseudovariety corresponding to piecewise k-testable languages for various k. It is interesting to observe that these semigroups all lie in the pseudovariety $\mathbb{L}\mathbf{DS}$, thus the direct product of any one of these with $\mathbf{B}_2^1$ generates a pseudovariety that lies in the interval $[\![\mathbb{V}(\mathbf{B}_2^1), \mathbb{L}\mathbf{DS} \vee \mathbb{V}(\mathbf{B}_2^1)]\!]$ so has NP-hard membership problem. It would be very interesting if this pseudovariety also has co-NP-hard membership problem (with respect to many-one reductions).

8. Variety membership for some small algebraic structures

An early result due to Lyndon [30] is that every two element algebra of finite signature has a finite basis for its equations. Because membership in the variety generated by such an algebraic structure is first order definable, its complexity lies in the complexity class $\mathtt{AC}^0$, a proper subclass of P. Thus the smallest possible size for an algebra to generate a variety with intractable membership problem is 3. To the author's knowledge, the smallest known algebra with intractable membership problem for its variety, assuming $\mathtt{P} \neq \mathtt{NP}$, has 6 elements: the original example due to Szekely [36] (subject to minor modification observed in Jackson and McKenzie [19, p. 123]). The NP-completeness results above now give an easy example of a 3-element algebra with NP-complete variety membership problem, as well as a 4-element groupoid with NP-hard variety membership. (Here by *groupoid* we mean

an algebraic structure consisting of a single binary operation.) Recall that the *graph algebra* of a graph $\mathbb{G}$ is the algebra formed over the vertices $V_{\mathbb{G}} \cup \{\infty\}$ (where ∞ is some new symbol) by letting ∞ act as a multiplicative zero element, and letting $u \cdot v = v$ if $\{u, v\} \in E_{\mathbb{G}}$ and $u \cdot v = \infty$ otherwise. Observe that for vertices $u_1, \ldots, u_n \in V_{\mathbb{G}}$ (not necessarily distinct), we have $(\ldots((u_1 \cdot u_2) \cdot u_3 \ldots) \cdot u_n \neq \infty$ if and only if $\{u_i, u_{i+1}\} \in E_{\mathbb{G}}$ for each $i \leq n-1$, in which case $(\ldots((u_1 \cdot u_2) \cdot u_3 \ldots) \cdot u_n = u_n$.

Corollary 8.1. *(1) It is NP-hard to decide membership of finite groupoids in the variety generated by the graph algebra of* $\mathbb{K}_3$ *(with 4-elements).*

(2) There is a 3-element algebra $\mathbf{M}$ *with two ternary operations such that it is NP-complete to decide membership of finite algebras in the variety generated by* $\mathbf{M}$*.*

Proof. (1). It is easy to see that if finite graph $\mathbb{G}$ lies in the universal Horn class of $\mathbb{K}_3$ then the graph algebra of $\mathbb{G}$ lies in the variety generated by the graph algebra of $\mathbb{K}_3$. Indeed, we may assume without loss of generality that $\mathbb{G}$ is an induced subgraph of a power $\mathbb{K}_3^\ell$ of $\mathbb{K}_3$. Let $\mathbf{A}$ denote the graph algebra of $\mathbb{K}_3$ and $\mathbf{B}$ denote the graph algebra of $\mathbb{G}$. Then $\mathbf{B}$ is a quotient of a subalgebra of $\mathbf{A}^\ell$: simply take the subalgebra of $\mathbf{A}^\ell$ generated by the vertices of $\mathbb{G}$ (which are ℓ-tuples of vertices in $\mathbb{K}_3$), and factor by the ideal consisting of all elements that have a coordinate equal ∞.

For the converse, assume that $\mathbb{G}$ is a finite connected graph *not* lying in the universal Horn class of $\mathbb{K}_3$. Let $u_1, \ldots, u_n$ be a sequence of vertices encountered in any path through $\mathbb{G}$ that includes each edge at least once in both directions (such paths obviously exist when $\mathbb{G}$ is connected, and moreover can be constructed in polynomial time, by taking any directed Eulerian circuit in the graph, treating the edge relation as a symmetric binary relation). Treat these vertex names as variables. Then the graph algebra of $\mathbb{K}_3$ satisfies $(\ldots((u_1u_2)u_3)\ldots u_n) \approx xx$, while the graph algebra of $\mathbb{G}$ fails this law under the trivial interpretation of variables as vertices. So the graph algebra of $\mathbb{G}$ does not lie in the variety generated by the graph algebra of $\mathbb{K}_3$. The result now follows from the NP-completeness of membership of connected graphs in the universal Horn class of $\mathbb{K}_3$ (which coincides with the NP-completeness of 2-robust 3-colourability of finite connected graphs).

(2). The idea is very similar, though now we can call directly on the construction of Jackson [18, Section 7.1], which produces a 3-element algebra $\mathrm{ps}(2^\pi)$ from the template 2 for monotone 1-in-3 3SAT. This algebra has one ternary operation (corresponding to the ternary relation of 2) and two binary operations $\wedge, \triangleright$. Proposition 7.2 of [18] then shows that the membership problem for finite algebras in the variety of $\mathrm{ps}(2^\pi)$ is polynomial time equivalent to the membership problem for the universal Horn class of 2, which is NP-complete by Theorem 6.1. The two binary operations $\wedge$ and $\triangleright$ can be replaced by the single ternary operation $(x \wedge y) \triangleright z$, due to satisfaction of the laws $(x \wedge y) \triangleright (x \wedge y) \approx x \wedge y$ and $(x \wedge x) \triangleright y \approx x \triangleright y$. □

References

1. J. Almeida, Finite Semigroups and Universal Algebra, World Scientific, Singapore 1994.
2. S. Abramsky, G. Gottlob, P.G. Kolaitis, Robust constraint satisfaction and local hidden variables in quantum mechanics, in IJCAI '13, pp. 440–446, 2013.
3. A. Beacham and J. Culberson, On the complexity of unfrozen problems, Disc. Appl. Math. 153 (2005), 3–24.

4. A.J. Beacham, The Complexity of Problems Without Backbones, MSc Thesis, University of Alberta, 2000.
5. C. Bergman and G. Slutzki, Complexity of some problems concerning varieties and quasivarieties of algebras, *SIAM J. Comput.* 30 (2000), 359–382.
6. A. A. Bulatov, P. G. Jeavons and A. Krokhin, *Classifying the complexity of constraints using finite algebras*, SIAM J. Comput. 34 (3) (2005), 720–742.
7. S. Burris and H.P. Sankappanavar, A Course in Universal Algebra, Graduate Texts in Mathematics 78, Springer Verlag, 1980.
8. J. Culberson and I. Gent, Frozen development in graph coloring, Theoret. Comput. Sci. 265 (2001), 227–264.
9. S. Eilenberg and M.P. Schützenberger, On pseudovarieties, Adv. Math. 19 (1976), 413–418.
10. L. Fortnow, J. Lutz, and E. Mayordomo, Inseparability and strong hypotheses for disjoint NP pairs. Theory Computing Systems, 51 (2012), 229–247.
11. C. Glasser, A.L. Selman, S. Sengupta and L. Zhang, Disjoint NP-pairs, SIAM J Comput. 33 (2004), 1369–1416.
12. O. Goldreich, On promise problems: a survey, in O. Goldreich et al. (Eds.): Shimon Even Festschrift, LNCS 3895, pp. 254–290, 2006.
13. V.A. Gorbunov, Algebraic Theory of Quasivarieties, Consultants Bureau, New York, 1998.
14. G. Gottlob, On minimal constraint networks, in J. Lee (Ed.): CP 2011, LNCS 6876, pp. 325–339, 2011.
15. J. Grollmann and A.L. Selman, Complexity measures for public key cryptosystems, SIAM J. Comput. 17 (1988), 309–335.
16. B. Hayes, Can't get no satisfaction, American Scientist 85 (1997), 108–112.
17. J.M. Howie, Fundamentals of Semigroup Theory, 2nd edition. Oxford University Press, New York (1995)
18. M. Jackson, Flat algebras and the translation of universal Horn logic to equational logic, Journal of Symbolic Logic 73 (2008), 90–128.
19. M. Jackson and R. McKenzie, Interpreting graph colourability in finite semigroups, Internat. J. Algebra Comput. 16 (2005), 119–140.
20. M. Jackson and G.F. McNulty, The equational complexity of Lyndon's algebra, Algebra Universalis 65 (2011), 243–262.
21. M. Jackson and B. Trotta, Constraint satisfaction, irredundant axiomatisability and continuous colouring, Studia Logica 101 (2013), 65–94.
22. M. Jackson and W.T. Zhang, From A to B to Z, manuscript, 2015.
23. O.G. Kharlampovich and M.V. Sapir, Algorithmic problems in varieties, Internat. J. Algebra Comput. 5 (1995), 379–602.
24. O. Klíma, Complexity issues regarding checking of identities in finite monoids, Semigroup Forum 79 (2009), 435–444.
25. O. Klíma, M. Kunc and L. Polák, Deciding piecewise k-testability, manuscript, 2014.
26. M. Kozik Computationally and algebraically complex finite algebra membership problems, Internat. J. Algebra Comput. 17 (2007), 1635–1666.
27. M. Kozik, A 2EXPTIME complete varietal membership problem. SIAM J. Comput. 38 (2009), 2443–2467.
28. B. Larose and P. Tesson, *Universal algebra and hardness results for Constraint Satisfaction Problems*, Theoret. Comput. Sci. 410 (2009), 1629–1647.
29. E.W.H. Lee and W.T. Zhang, Finite basis problem for semigroups of order six, LMS J. Comput. Math. 18 (2015), 1–129.
30. R. Lyndon, Identities in two-valued calculi. Trans. Amer. Math. Soc. 71 (1951), 457–465.
31. G.F. McNulty, Z. Szekely and R. Willard, Equational complexity of the finite algebra membership problem, Internat. J. Algebra Comput. 18 (2008), 1283–1319.
32. R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, L. Troyansky, Determining computational complexity from characteristic phase transitions, Nature 400 (1998), 133–137.
33. U. Montanari, Networks of constraints: Fundamental properties and applications to picture processing, Information Sciences 7 (1974), 95–132.
34. J. Nešetřil and A. Pultr, On classes of relations and graphs determined by subobjects and factorobjects, Disc. Math. 22 (1978) 287–300.
35. M.M. Stronkowski, Quasi-equational bases for graphs of semigroups, monoids and groups, Semigroup Forum 82 (2011), 296–306.

36. Z. Szekely, Computational complexity of the finite algebra membership problem for varieties, Internat. J. Algebra Comput. 12 (2002), 811–823.
37. B. Trotta, Residual properties of simple graphs, Bull. Austral. Math. Soc. 82 (2010), 488–504.
38. M.V. Volkov, S.V. Gol′dberg and S.I. Kublanovskiĭ, A minimal infinitely based semigroup with a polynomially recognizable variety. (Russian) Fundam. Prikl. Mat. 16 (2010), 105–122; translation in J. Math. Sci. (N. Y.) 177 (2011), no. 6, 847–859.
39. W. Wheeler, The first order theory of n-colorable graphs, Trans. Amer. Math. Soc. 250 (1979), 289–310.
40. W.T. Zhang and J.F. Luo, A new example of a minimal non-finitely based semigroup, Bull. Aust. Math. Soc. 84 (2011), 484–491.

Department of Mathematics and Statistics, La Trobe University, Victoria 3086, Australia

E-mail address: M.G.Jackson@latrobe.edu.au